

Learning Adaptive Evolutionary Computation for Solving Multi-Objective Optimization Problems

Citation for published version (APA):

Coppens, R. H. M., Reijnen, R., Zhang, Y., Bliet, L., & Steenhuisen, B. (2022). Learning Adaptive Evolutionary Computation for Solving Multi-Objective Optimization Problems. *arXiv*, 2022, Article 2211.09719.
<https://doi.org/10.48550/arXiv.2211.09719>

DOI:

[10.48550/arXiv.2211.09719](https://doi.org/10.48550/arXiv.2211.09719)

Document status and date:

Published: 01/11/2022

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Learning Adaptive Evolutionary Computation for Solving Multi-Objective Optimization Problems

R.H.M. Coppens¹, R.V.J. Reijnen¹, Y. Zhang¹, L. Blik¹, B. Steenhuisen²

¹ Eindhoven University of Technology

² Nobleo Manufacturing

remcoppens@outlook.com, r.v.j.reijnen@tue.nl, yqzhang@tue.nl, l.blik@tue.nl, berend.steenhuisen@nobleo.nl

Abstract

Multi-objective evolutionary algorithms (MOEAs) are widely used to solve multi-objective optimization problems. The algorithms rely on setting appropriate parameters to find good solutions. However, this parameter tuning could be very computationally expensive in solving non-trivial (combinatorial) optimization problems. This paper proposes a framework that integrates MOEAs with adaptive parameter control using Deep Reinforcement Learning (DRL). The DRL policy is trained to adaptively set the values that dictate the intensity and probability of mutation for solutions during optimization. We test the proposed approach with a simple benchmark problem and a real-world, complex warehouse design and control problem. The experimental results demonstrate the advantages of our method in terms of solution quality and computation time to reach good solutions. In addition, we show the learned policy is transferable, i.e., the policy trained on a simple benchmark problem can be directly applied to solve the complex warehouse optimization problem, effectively, without the need for retraining.

Introduction

In Combinatorial Optimization Problems (COPs), the goal is to identify high-quality solutions. Given that most COPs are NP-hard, solution approaches typically rely on heuristics to achieve good solutions in short computation time, compromising the optimality. Machine learning (ML) has recently been successfully leveraged for learning to solve classical COPs such as Travelling Salesman Problem (TSP) and other complex optimization problems (Kool, van Hoof, and Welling 2018; Kwon et al. 2020; Hottung, Bhandari, and Tierney 2020; da Costa et al. 2021).

In comparison to the amount of attention spend on single-objective COPs, much less work has been done in the ML community on solving multi-objective optimization problems (MOOPs), which are highly relevant problems in practice (Tian et al. 2021). In the optimization community, Multi-Objective Evolutionary Algorithms (MOEAs) are widely developed to solve MOOPs, through an iterative process of proposing solutions to the problem based on a set of previously proposed solutions. To optimize the performance of a MOEA, one needs to tune a set of parameter values. The process of tuning these parameters is often cumbersome and will lead to an (sub-)optimal set of values that is specifically tailored to the underlying problem. According to the “No

free lunch” theorem (Wolpert and Macready 1997), there does not exist a set of parameter settings that are superior to other settings in all stages of optimization, or on all optimization problems. Besides the challenge of finding generalizable parameter settings in MOEA or evolutionary algorithms in general, another known challenge is their high computation time in converging to good solutions, especially for solving complex and high dimensional optimization problems.

A way to circumvent the need to finetune the parameters of the MOEA is to implement Adaptive Parameter Control (APC) (Aleti and Moser 2016). By adaptively setting the values dictating the behavior of the operators within the algorithm, APC is able to tailor the different operator settings to the different stages of optimization. Besides allowing for increased performance, it also eliminates the cumbersome task executing parameter optimization. As a MOEA operates using an iterative process of creating a population based upon the previous population, sequential decision making shows beneficial for this task. For this reason Deep Reinforcement Learning (DRL) shows most salient, as it considers the optimization performance attained by the algorithm until the current generation while it proposes new parameter settings.

In this paper, we implement APC using DRL on a MOEA. More specifically, we develop a framework that uses a Dueling Deep Q-Network (DDQN) agent to dynamically set the operators dictating the behavior of the Non-dominated Sorting Genetic Algorithm III (NSGA-III). The performance of the framework is tested on both a simple benchmark problem and a real-world Warehouse Design and Control Problem (WDCP), which concerns simultaneous optimization of warehouse layout design and the underlying control policies. By simultaneous optimization of its constituent sub-problems, solving the WDCP allows to take the (hierarchical) interrelations between the problems into account. This results in higher solution quality, albeit at an increased computational cost. Furthermore, we demonstrate the generalisability of the proposed approach by solving WDCP with the policies trained on a much simpler optimization problem, without the need for adjustments to the underlying method and without retraining of the DDQN agent.

The contributions of this work are summarized as follows:

- We propose a framework for adaptive parameter control of a Multi-Objective Evolutionary Algorithm that generalizes to different problems. To the best of our knowledge,

we are the first to propose a DRL parameter control approach for a Multi-Objective Evolutionary Algorithm.

- We effectively train a Deep Reinforcement Learning policy for adaptive parameter control on a cheap-to-evaluate multi-objective optimization problem and deploy it to a real-world, computationally expensive, multi-objective combinatorial optimization problem.

Preliminaries and related work

Multi-Objective Optimization (MOO) is a branch of problems involving two or more objective functions to be simultaneously optimized. MOO can be mathematically expressed as follows:

$$\begin{aligned} & \text{Minimize } F(x) = \{f_1(x), f_2(x), \dots, f_m(x)\} \\ & \text{Subject to } g(x) \leq 0 \end{aligned}$$

where x is the vector of design variables, $f_i(x)$ is the i th objective function, and $g(x)$ is the constraint vector. The solution of MOO is generally expressed as a set of Pareto optima, representing the optimal trade-offs between the different objective values. Solutions present in this Pareto optima are Pareto optimum or non-dominated solutions, and a feasible solution x^* is Pareto optimal if there exists no feasible vector x such that $f_i(x) \leq f_i(x^*)$, $i \in \{1, 2, \dots, m\}$ and $f_i(x) < f_i(x^*)$, for at least one objective $i \in \{1, 2, \dots, m\}$.

The multi-objective evolutionary algorithm (MOEA) has been demonstrated as an effective method for solving MOO problems, among which the Non-dominated Sorting Genetic Algorithm III (NSGA-III) is one of the most successful ones (Deb and Jain 2013). MOEAs iteratively create a set of solutions, named a population of individuals, based on the well-performing solutions found in the previous population. The evolutionary process consists of three operators, the crossover operator, mutation operator and selection operator. The crossover operator takes two individuals, called parents, to create two new individuals, called childs, through recombination of the constituent parts of both parents. Thereafter the childs are taken through mutation, which randomly changes values to insert new values into the process. Finally, the parents and their childs are taken through the selection operator, which decides which individuals to remain and which to discard. Every cycle through all operators is called a generation, which are repeated until a set termination criterion is reached.

The DTLZ benchmark test suite is a widespread test suite for multi-objective problems with scalable fitness dimensions (Cheng et al. 2017). All problems in this test suite are continuous n -dimensional multi-objective problems.

The Hypervolume Indicator (HV) is a frequently used performance measure for multi-objective optimization (Zitzler and Künzli 2004). The hypervolume is described as the volume of the space in the objective space dominated by the Pareto front approximation Y_N , which is the set of optimal solutions found, and bounded from above by a reference point $r \in \mathbb{R}^m$ such that for all solutions $y \in Y_N$, $y \leq r$. The hypervolume indicator is computed as:

$$HV(Y_N; r) = \lambda_m \left(\bigcup_{y \in Y_N} [y, r] \right),$$

in which λ_m refers to the m -dimensional Lebesgue measure for assigning a measure to the m -dimensional Euclidean space. It is defined by first calculating the Lebesgue outer measure, which implies the infimum of all values on the interval comprised by the values in the Pareto frontier. The values obtained through calculation of this infimum, show the dominated area and thus the hypervolume indicator value. A larger value of the hypervolume indicates a better found approximated Pareto front. By tracking the hypervolume indicator over consecutive generations, the convergence rate of the Pareto front can be analyzed.

Deep reinforcement learning for parameter control.

Several existing works use DRL as APC in evolutionary computation to solve optimization problems. All these works consider the Differential Evolution (DE) algorithm. Sharma et al. (2019) uses a Double Deep Q-Network (DQN) agent to aid the selection of different mutation operators within DE. The state representation focused on the locations and distribution of separate individuals in the solution space. Tan and Li (2021) propose a somewhat similar approach using DQN, but their state representation focuses on the complexity of the solution space itself. Sun et al. (2021) implement a Policy Gradient method as the DRL algorithm. All these works use the CEC2013 and CEC2017 benchmarks to test the performance of their approaches. Our approach distinguishes itself in several aspects. First, we adopt the NSGA-III algorithm. The benefit of NSGA-III is the ability to maintain individual objective values throughout the process of multi-objective optimization. Doing this increases the utility of the approach for human decision-makers as they gain insight into the distribution of performance over the separate objective values. More importantly, with DRL as APC for NSG-III, the learned policy is highly generalizable. The agent is trained on a simpler problem with fast training time, and then it can be deployed to solve a much more complex problem. Through applying DRL as APC on the NSGA-III, we were able tackle a challenging, real-world optimization problem.

Solution framework DRL-MOEA

The intended purpose of applying DRL is to increase the convergence speed without damaging the performance of the NSGA-III. Doing this will keep the benefit of using NSGA-III, namely the optimization performance, while mitigating the downside, namely the high computational cost. The benefit of using the proposed framework, instead of the regular NSGA-III, becomes especially apparent when the underlying problem increases in complexity. The additional computational cost of training an agent can be mitigated by training the agent on a more simplistic optimization problem before deploying it to a more complex problem.

DRL for APC

We model the Adaptive Parameter Control (APC) problem for NSGA-III as a sequential decision problem. In the implementation of DRL in APC, it concerns learning policies of setting the parameter values dynamically during the iterations of the genetic algorithms to optimize the total reward.

The implementation of DRL will be done using the offline, model-free Deep Q-Network (DQN). DQN learns to predict the values of being in a certain state, and in every state selects the action leading to the next state with the highest value. Taking these most “profitable” actions will result in the creation of a so-called optimal behavioral policy, which tells the agent to take which action in which state. As evaluation time is the main concern of this research, training needs to happen using as little evaluations as possible. In this case, DQN has some beneficial characteristics. These characteristics concern sample efficiency and robustness, which are attained through implementation of experience replay and a frozen target network respectively. Experience replay concerns a buffer of interactions, frequently sampled to train the weights of the underlying neural network. The frozen target network, on the other hand, implies the agent using two neural networks. One network is used to dictate the behavior of the agent, where the other is used to predict the value of being in a certain state. The weights of the value network are repeatedly updated, using samples taken from the experience replay buffer. Only periodically, the learned weights will be copied onto the behavioral network, which dictates the behavior of the agent. This way the behavior of the agent remains stable, preventing it to fall into a positive feedback loop in which it chases an action returning a desired return.

The variation of the DQN used is called the Dueling DQN, or DDQN. The additional characteristic concerns a so-called Advantage function (Sewak 2019). This function returns an advantage value for all possible actions the agent can take, when in a certain state. These advantage values tell how much better it would be to take action a_k in state s over all other possible actions $a \in \mathcal{A}$ in state s . The advantage values together with the state values are used to calculate the Q-value(s), with $Q(s; a; \theta, \alpha, \beta) = V(s; \theta, \beta) + (A(s, a; \theta, \alpha) - \frac{1}{|\mathcal{A}|} \sum_{a'} A(s; a'; \theta, \alpha))$. The main benefit of using the advantage function besides the state values is to generalize learning across actions without imposing any change to the underlying reinforcement learning algorithm (Wang et al. 2016). Dueling DQN (DDQN) is beneficial if there could be cases where two actions have identical value, which is likely to happen given the environmental design the agent is deployed in.

State representation. The state representation concerns a description of the current state the agent is in. This representation will be taken through the behavioral network of the DDQN agent, which will return an action to be executed. Since the purpose of the proposed solution method is to be generalizable, all variables comprising the representation need to be made generic. The constructed representation focuses on three parts, namely, (1) the progress of the evolutionary process; (2) the performance and spread of the individuals comprising the population; and (3) the performance and size of the found Pareto optimal set.

The evolutionary progress is measured using two variables, the current generation (\mathcal{G}) and the stagnation counter (\mathcal{S}). The current generations (\mathcal{G}) gives the agent insight into how much time remains to optimize, where the stagnation counter (\mathcal{S}) implies how many generations it did not produce any

improvement. The stagnation counter is clipped, meaning that a maximal value is set. The stagnation counter value cannot exceed 10, which is a boundary that is only reached in extreme cases concluded from empirical evaluations.

The population performance and spread is summarized in three values. The first consists of the average of all normalized objective values (O_{mean}), providing an indication of the overall performance of the population. The second value concerns the average of all normalized minimal values per objective value (O_{min}), focusing on the best obtained values found so far. The third value describes the average normalized standard deviations (σ) of all objectives. The reason the values are aggregated into a single value is to allow for generalization. By using aggregates, the agent can be applied to any optimization problem, independent of the number of objective values considered.

The last part of the state representation describes the performance of the current approximation of the Pareto front. This performance is summarized in two metric, the hypervolume indicator (\mathcal{H}) and the size of the Pareto set (PS). The hypervolume indicator (\mathcal{H}) is the most used set-quality indicators for the assessment of multi-objective optimizers when the actual Pareto front is unknown (Guerreiro, Fonseca, and Paquete 2020). The Pareto size (PS) is included to show the agent to what extend the maximum size of the front is reached. If the pareto size has not reached its maximum, there is a possibility to extend the Pareto optimal set without the need to discard another solution.

Action space. The action space of the agent concerns the decisions it is able to make. With regard to NSGA-III, three values are considered to be included in this space. These values are summarized in 1 below, also showing the current values used in the initial implementation of the NSGA-III (Deb and Jain 2013) and describing the effect of high and low values.

The values described in 1 are solemnly considered. Actual construction of the action space relies on additional information gained from the process of Bayesian hyperparameter tuning, from which regions of well-performing hyperparameter settings can be extracted. The actual action space will be made discrete, having their values uniformly taken from the parametric ranges gathered from the hyperparameter tuning. The main reason the action space is made discrete is that small changes in the considered values do not impact behavior significantly. By setting discrete value pairs the selection of different action can be made significantly different, which we assume to be desirable for the implementation of the DDQN agent.

Reward function. The intended purpose of implementing DRL as APC in the NSGA-III is to improve its convergence speed while maintaining solution quality. Hence, the implementation of the DDQN agent is not immediately on the problem, instead its environment consists of the underlying evolutionary algorithm. For this reason, the objective of the agent is different than the actual objectives of the given optimization problem. The DRL objective concerns fast convergence to a Pareto front while maintaining the quality of the found set of solutions. Scale sensitivity is of high

Table 1: Action space Deep Reinforcement Learning (DRL) agent

	η_{SBX}	η_{PLM}	$indpb_{PLM}$
Operator	Crossover	Mutation	Mutation
Description	Crossover distribution parameter	Mutation distribution parameter	(Independent) mutation probability
NSGA-III value	30	20	0.01
Effect of high value	Produce children resembling their parents	Produce mutation resembling its origin	Higher chance for values to be mutated
Effect of low value	Produce children dissimilar to their parents	Produce mutation dissimilar to its origin	Lower chance for values to be mutated

importance when constructing a reward function for this purpose (Karafotias, Hoogendoorn, and Eiben 2015). Within the NSGA-III, this implies the relative difference in attainable improvement comparing the first generations to the final generations. Inspired by Huang et al. (2021), we define our reward function as an episodic reward, calculated by the sum of all hypervolume indicator values over an entire optimization run: $Reward_{gen} = \sum_{g=0}^{[1, \dots, \mathcal{G}]} HV_g, \forall gen \in [1, \dots, \mathcal{G}]$. This metric includes both convergence speed and convergence quality, as an increase in either will increase the metric value. Using this metric will result in the agent optimizing its behavior towards converging in as few generations as possible to a set of solution having the highest solution quality possible.

Warehouse Design and Control Problem (WDCP)

As stated earlier, the purpose of applying DRL is to increase the convergence speed without damaging the performance of the NSGA-III. Doing this will keep the benefit of using NSGA-III, namely the optimization performance, while mitigating the downside, being the high computational cost. The benefit of using the proposed framework, instead of the regular NSGA-III, becomes especially apparent when the underlying problem increases in complexity. To validate this statement, the proposed approach will first be trained and tested on a simple problem after which it will be deployed on a complex real-world problem, being the Warehouse Design and Control Problem (WDCP), without retraining. For the latter a simulation model is created, based on an actual warehouse of a large plastic manufacturer, to evaluate the performance of the individuals proposed by the NSGA-III. A visual representation of the entire solution framework for optimizing the WDCP, including the warehouse simulation, is shown in Figure 1.

The studied WDCP consists of three sub-problems, namely the warehouse lay-out design problem, the resource allocation problem and the product allocation problem. The decision variables can be subdivided over four categories. The first two categories concern values dictating the Product Placement Algorithm (PPA). These values concern four ordinal values (\mathcal{O}), dictating the order of the rules within the PPA, and five parameter values ($\delta_1, \delta_2, \delta_3, \delta_4, \delta_5$), dictating the actual behavior of every single rule. The third category of decision variables concerns resource allocation and decides the number of resources per resource type ($\mathcal{Z}_1 \dots \mathcal{Z}_n$). Finally, the fourth category concerns the dimensions of storage locations. These dimensions can differ based on either

width (\mathcal{W}), height (\mathcal{H}) or both. A formal definition of the number of storage locations having a given width/height combination is however trivial. An abstract representation of the amount of storage locations per hall (\mathcal{X}) can be described as $d_{w,h}^x \forall x \in \mathcal{X}, \forall w \in \mathcal{W}, \forall h \in \mathcal{H}$.

The WDCP concerns multiple objectives: the tardiness of outbound trucks, the total cost of resources, and the number of unplaceable products. Upon arrival of an inbound truck (\mathcal{K}_{inb}) the truck is placed at one of the truck docks (o_k). After docking, all products are retrieved and placed throughout the warehouse in the locations determined using the PPA, dictated by the rules order (\mathcal{O}) and the parameter values ($\delta_1, \delta_2, \delta_3, \delta_4, \delta_5$). Transportation to these storage locations is executed using the available resources ($\mathcal{Z}_1 \dots \mathcal{Z}_n$). The number of storage locations of different dimensions in the different storage halls (\mathcal{X}) depends on $d_{w,h}^x \forall x \in \mathcal{X}, \forall w \in \mathcal{W}, \forall h \in \mathcal{H}$. As products can only be stored in storage areas equal or larger in either one or both of the dimensions, the possible storage locations of a given storage type (\mathcal{S}) for a given product can be defined as $d_p^{\mathcal{S},x} \forall w \in [p_w, \dots, \mathcal{W}^{\mathcal{S}}], \forall h \in [p_h, \dots, \mathcal{H}^{\mathcal{S}}], \forall x \in [1, \dots, \mathcal{X}^{\mathcal{S}}]$. If no available storage location can be found, the product is classified unplaceable, removed from the simulation and the average transportation time is used for the time a resource is occupied for placing the product in the warehouse.

The remaining two objective values concern the outbound process. Upon arrival of an outbound truck (\mathcal{K}_{outb}), it is placed at one of the available truck docks (o_k). Thereafter the requested products are retrieved one-by-one using the available resources ($\mathcal{Z}_1 \dots \mathcal{Z}_n$). The tardiness of these outbound trucks (y_k) is calculated by taking the difference between the departure time (d_k) and arrival time (a_k), multiplied with a penalty factor (f_k). This penalty takes a value of 0 for trucks departing within 30 minutes, 0.5 if trucks depart between 30 and 120 minutes and 1 for truck having a departure time exceeding 120 minutes. A mathematical formulation for this objective value is $y_k = \sum_{k \in \mathcal{K}_{outb}} f_k \cdot y_k$. The final objective value, total resource cost, is calculated as $\sum_{i=1, \dots, n} c_i \cdot \mathcal{Z}_i$, where \mathcal{Z}_i describes the amount of resources for resource type i and c_i concerns an aggregate of both investment and operating costs of that resources.

Optimization is executed using the NSGA-III algorithm. As previously explained, the NSGA-III proposes solution methods based upon previously evaluated solutions. A proposed solution consists of 96 values, including the rule order values (\mathcal{O}), the rule parameter values ($\alpha, \beta, \gamma, \delta, v$), the

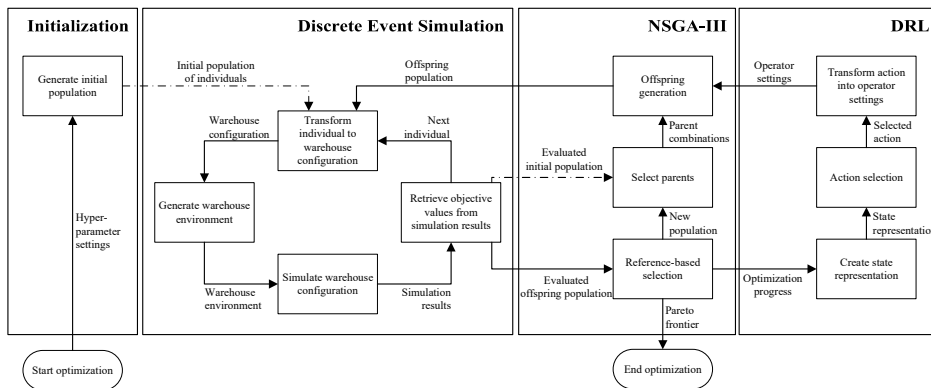


Figure 1: DRL-MOEA for WDCP.

amount of resources per resource type ($Z_1 \dots Z_n$) and the amount of storage locations dedicated to a given set of dimensions per storage hall ($d_{w,h}^x \forall x \in \mathcal{X}, \forall w \in \mathcal{W}, \forall h \in \mathcal{H}$). These solutions are given to the Discrete Event Simulation (DES), which creates an environment based upon the given values. Thereafter the configuration is evaluated by running a simulation of 48 hours, of which 8 hours is a warm-up period to prevent a simulation bias due to starting with a warehouse that is not in equilibrium yet.

The NSGA-III repeats the generational cycle by continuously proposing solutions, combining and mutation solutions and proposing new ones. Within every generation the framework creates a state representation right after the selection process of the NSGA-III. This state representation is thereafter used by the DDQN agent to dynamically set the values dictating the behavior of the mutation operator. Through dynamic operator control, the algorithm gets tailored to the progress made in optimization so far. This process is repeated until a given termination criterion is reached, which is set to 200 generations.

Experimental results

All experiments were run on a single computer containing 12 Intel I7-8750H CPUs, having 15.5 GB RAM at its disposal. The results regarding optimization performance show both an average and standard deviation of a set number of independent optimization trajectories. The amount of independent optimization runs used is dependent on the computation cost incurred and varies between the simple and complex problem. We compare the proposed DRL-MOEA approach with (1) the NSGA-III with optimized hyperparameter settings (optimized NSGA-III), and (2) an untrained (random) implementation of DRL-MOEA, which is the NSGA-III with adaptive parameter control (APC) where parameters are randomly selected, in contrast to those being learned in DRL-MOEA. The parameter settings used for all experiments can be found in the provided Supplementary Material.

Performance on the DTLZ2 benchmark

First, the agent is trained and tested using simple problems from the DTLZ benchmark test suite. Since results are sim-

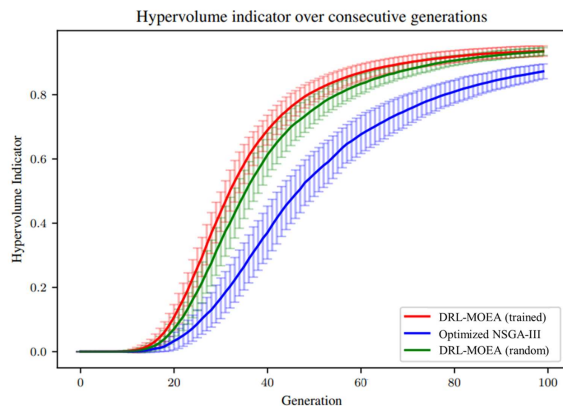


Figure 2: Performance benchmark DRL agent on the DTLZ2 benchmark

ilar, we report here only the performance on the DTLZ2 benchmark function. After 4000 episodes of 200 generations each, the performance of the learned policy is evaluated. Due to the relatively low computation cost, 500 independent optimization trajectories are used to evaluate the performance of the different models. To emphasize the effect on convergence speed, the evaluation of performance will focus on the first 100 generations. The remaining 100 generations will show an almost linear line, indicating the convergence to an optimal set of solutions to the problem. The attained performances of the different models is shown in 2.

Interesting to see is that both methods that apply APC, namely DRL-MOEA (trained) and DRL-MOEA (random), show an increase in optimization performance with regards to convergence speed. A possible explanation for the increased performance of the DRL-MOEA (random) can be found in both the use of a small population size of 20 individuals and the structural design of the NSGA-III. Using a small population size applying radical mutation can be beneficial to escape local optima. Under the expectation that the DRL-MOEA (random) will decide uniformly which action to take, 50% of its actions will show an inclination towards radical mutation. Additionally, the selection operator within the NSGA-III di-

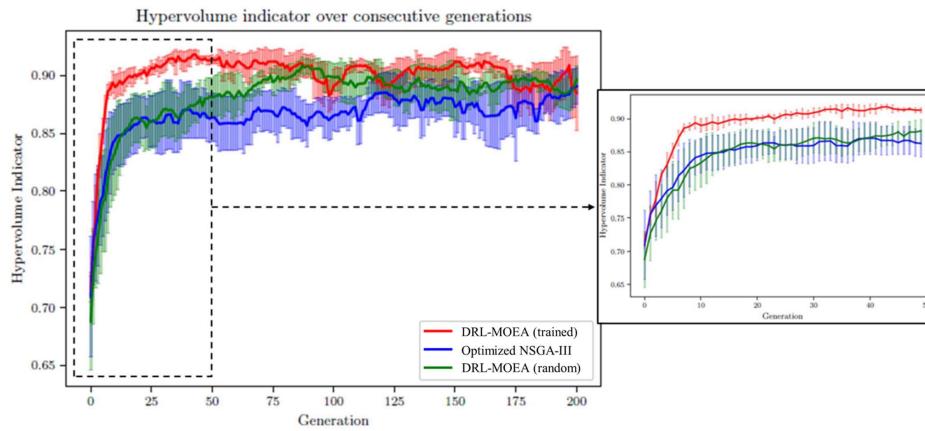


Figure 3: Performance benchmark DRL agent on the WDCP

minishes the downside of radical mutation, as individuals with poor performance will be removed from the population. But still, the learned DRL-MOEA outperforms the random DRL-MOEA, which indicates the added benefit of learning a policy to apply mutation within the NSGA-III.

Apply learned policy to solve WDCP

After learning an optimal policy on the DTLZ2 benchmark, the agent will be applied on the more complex problem, namely the Warehouse Design and Control Problem (WDCP). Solving the WDCP is done on an actual real-world use case, for which the warehouse of a large plastic manufacturer is replicated. Due to the high dimensionality and high level of detail of this simulation, the computation cost for a single evaluation is around 2-4 seconds. For this reason only five independent optimization trajectories are used to evaluate the three models. The results obtained are visualized in 3.

Figure 3 shows the performance of the learned agent (i.e. DRL-MOEA (trained)) is superior to the performance of both the DRL-MOEA (random) and optimized NSGA-III, even though the agent was trained on a different problem, namely DTLZ2. The learned DRL-MOEA (random) attained a maximum hypervolume indicator of 0.92, in comparison to 0.91 of the random DRL-MOEA and 0.89 of the optimized NSGA-III. Until generation 50 the increased performance is significant, as emphasized in the cut-out. After generation 50, the NSGA-III with a random action selecting agent implementation of APC shows to gain performance, which might indicate it breaking free from a local optimum due to the previously mentioned probability of an inclination towards radical mutation. Besides convergence speed, also the volatility over the five consecutive optimization trajectories decreased using the learned agent.

To further evaluate the robustness of the learned agent, an additional sensitivity analysis is conducted. In this analysis three scenarios are tested, which all focus on one or more of the constituent sub-problems of the WDCP. The first analysis increases the amount and size of incoming trucks, increasing overall pressure on the warehouse. The second analysis removes the consistency of truck arrivals, increasing the com-

plexity of the resource allocation problem. The third analysis increases the product portfolio, doubling the amount of distinct products, resulting in higher complexity of the product allocation problem. The results of all three analysis are shown in Fig. 4, showing the difference in performance of using a learned agent or not. As with the initial analysis, the performances are calculated over five independent optimization trajectories. In two of the three analysis, the learned agent as APC in the NSGA-III has led to increased performance. The maximum obtained hypervolume indicator for the learned agent and the optimized NSGA-III without an agent are 0.79 and 0.77 for the first experiment, 0.86 and 0.86 for the second experiment and 0.92 and 0.90 for the third experiment respectively. For the second experiment it is expected that the performance attained by the optimized NSGA-III without APC is already the optimal performance attainable, thus no improvement was possible by implementing APC using the learned agent. Besides, the volatility between the five independent optimization trajectories decreased by using the learned agent in all three scenarios, indicating a more stable optimization process.

Policy analysis

Both on the DTLZ2 benchmark as well as the WDCP the learned agent outperforms the benchmarked methods. The fact that the learned agent is able to retain superior performance shows that the learned behavior is generalizable, as no additional training is executed for the WDCP. To gain further insight into the behavior of the learned agent, the action selection policy for both the DTLZ2 benchmark as well as the WDCP is further analyzed. The action space of the agent consists of 15 different combinations of mutation intensity (η) and probability (indpb), where lower values indicate conservative mutation and higher values indicate more radical mutation.

Besides gaining insight into the action selection behavior of the agent, this analysis can also be used to validate the need to use DRL in the first place. If the behavior for both problems shows to be identical, the agent possibly can be replaced by a set of rules deduced from the agent. If the




Figure 4: Sensitivity analysis of the learned agent on the WDCP

behavior shows significantly different, it indicates the need for sequential decision and thus to use DRL.

Analysis of the behavioral policy for the DTLZ2 benchmark is based upon the actions taken in all 500 independent optimization trajectories. The fraction of action selection in every generation is shown in 5. The agent behavior indicates that the agent starts with more conservative mutation until generation 20, after which the inclination towards more radical mutation increases between generation 20 and 50. The initial preference for conservative mutation is however counter intuitive. One would expect that such algorithms tend to first explore the search space after which they exploit the gained knowledge in exploration to improve. The behavior shown in 5 indicates the opposite, as the agent first inclines towards exploitation after which it gradually moves towards more exploratory behavior.




Figure 5: Policy visualization for the DTLZ2 benchmark problem over consecutive generations

We analyze the behavior of the learned agent on the WDCP. We ran 40 independent optimization trajectories using the learned agent. The decrease in trajectories is due to the increased computation cost as a result of the simulation used in the WDCP. Figure 6 shows the fraction of actions selected over consecutive generation for 40 independent optimization trajectories. Concluding from the action selection behavior of the learned agent in Figure 6 an identical shift from conservative to radical mutation can be seen. However, the increased complexity of the underlying problem resulted in the agent being more decisive in its action selection. Compared to the DTLZ2 benchmark, in which all actions were selected at least




Figure 6: Policy visualization for the WDCP over consecutive generations)

once, only five different actions are selected by the agent on the WDCP. Also, the agent tends toward a higher intensity of mutation as opposed to the DTLZ2 behavior. This might be due to the WDCP being a non-convex optimization problem, which increases the need for radical mutation to escape the local optima in the search space.

Conclusions

Through this research, a generalizable solution framework is proposed that integrates the NSGA-III with APC, implemented using a Dueling DQN agent. The proposed framework initially shows increased performance on the DTLZ2 benchmark, which is the problem used throughout training of the agent. Thereafter the framework is applied to a complex real-world problem. The performance attained by the framework shows to be superior to that of the optimized NSGA-III without APC. The fact that the performance was retained when the underlying problem was changed and the model was not retrained, shows that the model is generalizable. Through a sensitivity analysis, increasing the complexity of the constituent sub-problems of the WDCP, this performance increase is maintained, which indicates the robustness of the framework. A difference in behavioral policy of the agent between the DTLZ2 benchmark and the WDCP shows the inability to replace the DRL agent with a simplistic set of rules. This validates the need to use DRL as a sequential decision making framework.

References

- Aleti, A.; and Moser, I. 2016. A systematic literature review of adaptive parameter control methods for evolutionary algorithms. *ACM Computing Surveys (CSUR)*, 49(3): 1–35.
- Cheng, R.; Li, M.; Tian, Y.; Zhang, X.; Yang, S.; Jin, Y.; and Yao, X. 2017. A benchmark test suite for evolutionary many-objective optimization. *Complex & Intelligent Systems*, 3(1): 67–81.
- da Costa, P.; Rhuggenaath, J.; Zhang, Y.; Akcay, A.; and Kaymak, U. 2021. Learning 2-Opt Heuristics for Routing Problems via Deep Reinforcement Learning. *SN Computer Science*, 2(5): 1–16.
- Deb, K.; and Jain, H. 2013. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints. *IEEE transactions on evolutionary computation*, 18(4): 577–601.
- Guerreiro, A. P.; Fonseca, C. M.; and Paquete, L. 2020. The hypervolume indicator: Problems and algorithms. *arXiv preprint arXiv:2005.00515*.
- Hottung, A.; Bhandari, B.; and Tierney, K. 2020. Learning a latent search space for routing problems using variational autoencoders. In *International Conference on Learning Representations*.
- Huang, C.; Li, L.; He, C.; Cheng, R.; and Yao, X. 2021. Operator-Adapted Evolutionary Large-Scale Multiobjective Optimization for Voltage Transformer Ratio Error Estimation. In *International Conference on Evolutionary Multi-Criterion Optimization*, 672–683. Springer.
- Karafotias, G.; Hoogendoorn, M.; and Eiben, A. 2015. Evaluating reward definitions for parameter control. In *European Conference on the Applications of Evolutionary Computation*, 667–680. Springer.
- Kool, W.; van Hoof, H.; and Welling, M. 2018. Attention, Learn to Solve Routing Problems! In *International Conference on Learning Representations*.
- Kwon, Y.-D.; Choo, J.; Kim, B.; Yoon, I.; Gwon, Y.; and Min, S. 2020. Pomo: Policy optimization with multiple optima for reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 21188–21198.
- Sewak, M. 2019. Deep q network (dqn), double dqn, and dueling dqn. In *Deep Reinforcement Learning*, 95–108. Springer.
- Sharma, M.; Komninos, A.; López-Ibáñez, M.; and Kazakov, D. 2019. Deep reinforcement learning based parameter control in differential evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 709–717.
- Sun, J.; Liu, X.; Bäck, T.; and Xu, Z. 2021. Learning adaptive differential evolution algorithm from optimization experiences by policy gradient. *IEEE Transactions on Evolutionary Computation*.
- Tan, Z.; and Li, K. 2021. Differential evolution with mixed mutation strategy based on deep reinforcement learning. *Applied Soft Computing*, 111: 107678.
- Tian, Y.; Si, L.; Zhang, X.; Cheng, R.; He, C.; Tan, K. C.; and Jin, Y. 2021. Evolutionary large-scale multi-objective optimization: A survey. *ACM Computing Surveys (CSUR)*, 54(8): 1–34.
- Wang, Z.; Schaul, T.; Hessel, M.; Hasselt, H.; Lanctot, M.; and Freitas, N. 2016. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, 1995–2003. PMLR.
- Wolpert, D. H.; and Macready, W. G. 1997. No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, 1(1): 67–82.
- Zitzler, E.; and Künzli, S. 2004. Indicator-based selection in multiobjective search. In *International conference on parallel problem solving from nature*, 832–842. Springer.