



HAL
open science

Parametric-Task MAP-Elites

Timothée Anne, Jean-Baptiste Mouret

► **To cite this version:**

Timothée Anne, Jean-Baptiste Mouret. Parametric-Task MAP-Elites. GECCO 2024, Jul 2024, Melbourne, Australia. 10.1145/3638529.3653993 . hal-04532964

HAL Id: hal-04532964

<https://hal.science/hal-04532964>

Submitted on 4 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Parametric-Task MAP-Elites

Timothée Anne

Jean-Baptiste Mouret

jean-baptiste.mouret@inria.fr

Université de Lorraine, CNRS, Inria

F-54000 Nancy, France

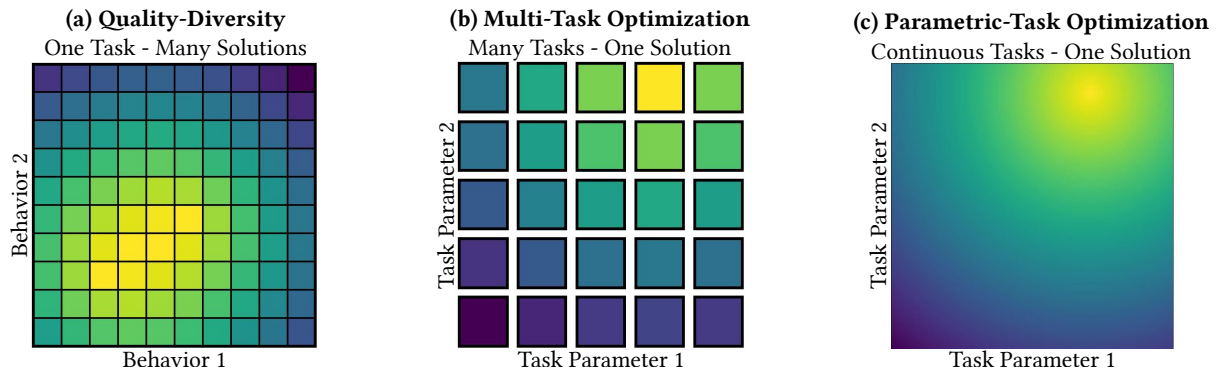


Figure 1: (a) **Quality-Diversity** is the problem of finding high-performing solutions with diverse behaviors. (b) **Multi-task optimization** is the problem of finding the optimum solutions for a finite set of tasks, each often characterized by a task parameter (or descriptor). (c) In this paper, we propose to extend the multi-task optimization problem to a continuous parametrization, which we call **Parametric-Task Optimization**. The goal is to be able to return the optimal solution for any task parameter.

ABSTRACT

Optimizing a set of functions simultaneously by leveraging their similarity is called multi-task optimization. Current black-box multi-task algorithms only solve a finite set of tasks, even when the tasks originate from a continuous space. In this paper, we introduce Parametric-Task MAP-Elites (PT-ME), a new black-box algorithm for continuous multi-task optimization problems. This algorithm (1) solves a new task at each iteration, effectively covering the continuous space, and (2) exploits a new variation operator based on local linear regression. The resulting dataset of solutions makes it possible to create a function that maps any task parameter to its optimal solution. We show that PT-ME outperforms all baselines, including the deep reinforcement learning algorithm PPO on two parametric-task toy problems and a robotic problem in simulation.

KEYWORDS

MAP-Elites, Multi-Task, Quality-Diversity, Robotics

ACM Reference Format:

Timothée Anne and Jean-Baptiste Mouret. 2024. Parametric-Task MAP-Elites. In *Proceedings of PREPRINT*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.48550/arXiv.2402.01275>

1 INTRODUCTION

Optimization is ubiquitous in today’s world [19]. While many problems require optimizing one cost function, many involve a set of related functions, each representing a different “task”. For example,

in robotics, it can be useful to learn walking gaits for various legged-robot morphologies [45], with each task corresponding to maximizing the walking speed for one morphology. Instead of solving each optimization individually, solving all the tasks simultaneously should decrease the total computation cost. Another illustration can be found in machine learning, where optimizing the hyperparameters of an algorithm for specific datasets can make a difference; here, each task is the loss function minimization for each dataset [32].

Multi-task problems are encountered in many fields such as robotics [7, 45], embodied intelligence [6, 45, 57], data science for symbolic regression of different real-dataset using genetic programming [62], multi-software testing using multifactorial evolutionary algorithm [52], logistics planning for different vehicle routing problems [17], complex design for cars [59], packaging [12], photovoltaic models [36], or electric power dispatch [40].

To our knowledge, all black-box multi-task algorithms solve a finite set of tasks, whereas many problem specifications are continuous. This means that continuous task spaces have to be discretized. For example, in robotics, the morphology is characterized by continuous length of limbs and joint bounds, and any discretization is arbitrary. When the cost function is differentiable, a few algorithms exist (i.e., Parametric programming [50]) to solve the continuous multi-task problem, which we call the parametric-task optimization problem, but we are not aware of any black-box algorithm to do so.

The contribution of this paper is a black-box parametric-task algorithm called Parametric-Task MAP-Elites (PT-ME) inspired by Multi-Task MAP-Elites (MT-ME) [45], an elitist multi-task algorithm. Our main idea is to design an algorithm that solves as many tasks as possible within its allocated budget, asymptotically solving

the parametric-task optimization problem by filling the task space with high-quality solutions. The resulting dense dataset can then be distilled, for instance, with a neural network, making it possible to generalize and obtain a high-quality solution for any task.

PT-ME’s key features are the following:

- it samples a new task at each iteration, meaning that the more evaluation budget it has, the more tasks it will solve;
- its selection pressure is de-correlated from the number of tasks it solves, meaning that its efficiency does not decrease with the number of tasks to solve;
- it uses a specially designed variation operator that exploits the multi-task problem structure to improve the performance;
- it distillates those solutions into a continuous function approximation, i.e., a neural network, effectively solving the parametric-task optimization problem.

We evaluate our method and its ablations extensively on two parametric-task optimization toy problems, 10-DoF Arm (used in [33, 45, 57]) and Archery (a more challenging problem with sparse rewards), and a more realistic parametric-task optimization, Door-Pulling, consisting of a humanoid robot pulling a door open. We compare against several baselines and show that our method effectively solves the parametric-task optimization problem while being more efficient than all baselines and ablations.

2 PROBLEM FORMULATION

The goal is to find a function G that outputs the optimal solution for each task parameter θ . More formally:

$$\forall \theta \in \Theta, G(\theta) = x_{\theta}^* = \underset{x \in \mathcal{X}}{\operatorname{argmax}}(f(x, \theta)) \quad (1)$$

where \mathcal{X} is the solution space, Θ the task parameter space, $f : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$ the function to optimize (also called fitness function), and $G : \Theta \rightarrow \mathcal{X}$ a function that outputs the optimal solution x_{θ}^* maximizing the fitness function for the task parameter θ . The main difference with multi-task optimization is that the task space Θ is continuous and not a finite set of tasks.

For ease of use in the remaining of the paper, we set the solution space $\mathcal{X} = [0, 1]^{d_x}$ where d_x is its dimension and the parameter space $\Theta = [0, 1]^{d_{\theta}}$ where d_{θ} is its dimension.

3 RELATED WORK

3.1 Parametric Programming

Optimizing a function parameterized with discrete or continuous parameters is studied in mathematics and is called parametric programming (or multiparametric programming for multiple parameters) [8, 18, 24, 49, 50]. The main algorithms divide the parameter space into critical regions and solve the optimization problem on each critical region using classical mathematical tools, i.e., Lagrange multipliers and Karush-Kuhn-Tucker conditions. There are three families of methods: multiparametric linear programming (mpLP) [24], multiparametric quadratic programming problems (mpQP) [49], and multiparametric nonlinear programming [18], which approximates nonlinear functions.

One application is to use parametric programming to precompute the optimal solutions for all parameter space and then query those solutions online. The reformulation of model predictive control

problems with a quadratic cost function into mpQP [50] has led to MPC on chip (or explicit MPC) with applications, for example, to control chemical plants [16].

All parametric programming algorithms assume that the function to optimize is known. While there are many situations where that information is available, for many real-world applications (e.g., complex simulations or real systems), the functions to optimize are non-linear, non-convex, non-differentiable, or even black-box, making those methods inapplicable.

3.2 Multi-Task Optimization

In recent years, multi-task optimization has seen a surge of interest in the evolutionary computation community [28, 29, 61]. We can decompose it into two families: implicit transfer of solutions (i.e., one solution space for all tasks) and explicit transfer of solutions (i.e., different solution spaces). Implicit transfer is more straightforward as the algorithm evolves only one population of solutions for all tasks but only applies to homogenous tasks. Explicit transfer concerns problems with heterogeneous tasks where knowing which information to share between tasks is not trivial. Proposed methods use probabilistic search distribution [27], search direction vectors [58], higher-order heuristics [31], or surrogate models [42].

Those methods have applications in many fields but rarely consider more than a handful of tasks simultaneously (i.e., between 2 and 10): 3 morphologies for robot controller using neuroevolution [6], 2 datasets for symbolic regression using genetic programming [62], 10 numerical calculus functions in C for software testing or 3 car models for design optimization using a multifactorial evolutionary algorithm [52, 59], 2 auxiliary tasks and one main task (i.e., 3 tasks) for safe multi-UAV path planning [5] or 2 bus systems power dispatch optimization using a multi-objective multifactorial evolutionary algorithm [40], 4 package delivery problems planning using explicit transfer knowledge [17], 3 diode models design using similarity-guided evolutionary multi-task optimization [36].

Other fields, such as Bayesian optimization [9, 54], propose multi-task algorithms, for example, to tune machine learning algorithms for several datasets [47]. Bayesian optimization is more focused on expensive fitness functions and cannot handle more than a thousand evaluations due to the cubic scaling of the Gaussian processes.

A handful of works tackle more than a dozen tasks: 30 in [38], 50 in [37], 500 in [57], and 2 000 in [33]. In this paper, we take inspiration from MT-ME [45], which scales up to 5 000 tasks.

3.3 Multi-Task MAP-Elites

Originating from novelty search [35], the Multi-dimensional Archive of Phenotypic Elites (MAP-Elites) algorithm [44] finds a large set of high-quality and diverse solutions to a task (this problem is called Quality-Diversity (QD)). MAP-Elites [44] has shown promising results in robotics with evolving a repertoire of diverse gaits for legged robots [11], or in video-gaming industries to generate procedural contents like spaceships and game levels [25]. Covariance Matrix Adaptation MAP-Annealing [21, 22] mixes MAP-Elites [44] with the self-adaptation techniques of CMA-ES [30] and outperforms both methods for finding diverse game strategies for the Hearthstone game. None of those variants have been applied to

multi-task optimization. This paper takes root from MT-ME [45], which is closer to the original MAP-Elites [44].

MT-ME [45] has applications in robotics, e.g., instead of optimizing each morphology individually, it simultaneously optimizes walking gaits for 2 000 morphologies. The resulting walking gaits allow robots to travel three times farther than those evolved with CMA-ES [30] run individually with the same budget of iterations. It has applications in industrial scenarios, such as optimizing ergonomic behaviors for different worker morphologies in an industrial workstation [63]. Multi-Task Multi-Behavior MAP-Elites [2] is an extension that solves the multi-task Quality-Diversity problem and outperforms MAP-Elites [44] run on each task individually.

MT-ME [45] explicitly discretizes the task space into a finite set of several thousand tasks and updates an archive of the best-known solution for each task, called an elite. At each iteration, it randomly selects two elites from the archive, generates a candidate solution using a variation operator, and evaluates this candidate solution on one of the tasks. If the task has no known solution or the candidate solution has a higher fitness than the current one, it becomes the elite for this task, and the algorithm resumes its main loop.

A first limitation of this discretization is that MT-ME [45] can only solve a finite number of tasks even though the original task space is continuous. Discretization does not scale well with dimensions, e.g., a thousand tasks only correspond to 32 steps per dimension in 2D and 10 steps per dimension in 3D. Another limitation is that, by being an elitist algorithm, its efficiency in finding solutions is correlated to the selection pressure, which is inversely correlated to the number of cells (i.e., the more cells there are to fill, the easier it is for a bad solution to become an elite, and the longer it takes to fill each cell with a good solution). As the number of cells equals the number of tasks to solve, the more tasks there are to solve, the slower the algorithm is to solve them.

3.4 Reinforcement Learning

As stated Sec. 2, the end product of parametric-task optimization algorithms should be a function $G : \Theta \rightarrow \mathcal{X}$ taking a task parameter θ as input and outputting the corresponding optimal solution x_{θ}^* . This function is similar to the action policy in Deep Reinforcement Learning (DRL) [3, 23], $\Pi : S \rightarrow A$ taking the current state descriptor $s \in S$ as input and outputting the corresponding optimal action $a \in A$ with regards to a long-term reward $r \in \mathbb{R}$. We can see the similarity between the task parameter θ and the state s , the solution x and the action a , and the fitness f and the reward r . The main difference is that parametric-task optimization algorithms optimize the fitness value directly. In contrast, DRL algorithms solve a more complex problem as they optimize the action for a long-term reward and must solve the credit assignment problem. Nevertheless, a DRL algorithm applied with a one-step horizon solves a parametric-task optimization problem. We make this parallel to highlight that, like DRL today, parametric-task optimization could have applications in many fields. Proximal Policy Optimization (PPO) [53] is a mature and widely used [10, 26, 34, 39, 48, 55, 60] DRL algorithm, that we think to be a strong baseline.

4 METHOD

Parametric-Task MAP-Elites (PT-ME) follows the same principles as MAP-Elites [44]. After initializing each archive cell with a random elite (Alg.1, Line.19), the main loop consists in generating a new solution candidate, evaluating it for a new task (L.41), and updating the archive by replacing the elite with the new solution if its fitness is greater or equal (L.43-46). The first difference is that we store all the evaluations (L.42) for later use during distillation (Sec.4.5) and evaluation (Sec.5.2.3). The second difference is that PT-ME uses MT-ME [44]’s variation operator for half of the iterations (Sec.4.2) and uses a new variation operator for the other half (Sec.4.3).

4.1 The Archive of Elites

Our first contribution is de-correlating the number of cells in the archive from the number of tasks solved. Like MT-ME [45], PT-ME divides the task parameter space into regions using CVT [15, 43] (L.13) and attributes an archive cell for each specific centroid task parameter θ_c (L.16-21). Nonetheless, instead of evaluating the candidate solutions only on those tasks, PT-ME samples a task parameter θ at each iteration (L.29-31 and L.35) and assigns it to the cell with the closest centroid θ_c (L.32 and L.36). By doing so, PT-ME improves the task parameter space coverage at each iteration and directly benefits from a larger budget of evaluations without reducing the convergence to good solutions in the early stage of the algorithm. One potential drawback (that we did not observe in our experiments) is that solutions for close but different tasks can be in competition because they are in the same cell, which means that an “easy” task could become the cell’s elite and prevent the algorithm from finding solutions to “harder” tasks of the same cell. This effect is easily attenuated by increasing the number of cells.

By contrast with MAP-Elites, the number of final solutions in PT-ME will far exceed the number of cells. However, the number of cells still impacts the selection pressure on the elites. At the extremes, with one cell, the selection pressure is maximal as every solution competes with every other, and with an infinite number of cells, the pressure is null as each solution is an elite. In addition, the number of cells impacts the new variation operator (Sec. 4.3). We found empirically that a value of $N = 200$ cells is robust for the parametric-task optimization problems presented Sec. 5.1.

4.2 Variation Operator 1: SBX with tournament

Like MT-ME [45] (but only for half of the iterations), PT-ME biases with a tournament the task selected to evaluate the candidate solution. However, instead of sampling the tasks among a fixed set, it uniformly samples them from the task parameter space.

More formally, for half of the iterations (L.28-33), PT-ME randomly selects two parents, p_1 and p_2 , from the archive of elites (L.28), samples s task parameters $\theta_{1:s}$ (L.29), selects the one closest to the task parameter from which p_1 was evaluated (L.31), and generates an offspring x with SBX [1] (L.33). The intuition is that two close tasks have more chances to share close solutions but that always taking the closest tasks would lead to poor exploration.

Like MT-ME [45], PT-ME chooses the size s of the tournament with a bandit (L.49-50), between a minimal size of 1 (i.e., no tournament) and maximal size of 500 (i.e., choosing a close task). PT-ME

also uses the UCB1 [4] bandit algorithm, which achieves the optimal regret up to a multiplication constant. The bandit score also corresponds to the candidate solution becoming an elite (L.47-48).

We concluded in a preliminary study that the bandit is significantly better than choosing the tournament size s at random but leads to slightly worse performances than fine-tuning it for each problem. We decided to keep it as it removes the need to fine-tune a hyperparameter without a significant loss in performance.

4.3 Variation Operator 2: Linear Regression

SBX [1] does not exploit the multi-task framework because it does not use the fact that we know the task on which the candidate solution will be evaluated. We take inspiration from the tournament, but instead of biasing the task parameter given the parents of the candidate solution, we bias the candidate solution given the task parameter. The idea is to aggregate information from the current archive by building a local model of the function $G : \theta \mapsto x_\theta^*$ and use it to estimate a candidate solution. We chose a linear model to “guess” the solution for the chosen task given the archive because it is simple and fast.

For half of the iterations (L.35-40), PT-ME:

- (1) samples a task parameter θ (L.35);
- (2) looks for the closest centroid θ_c (L.36) using a KDTree [41] with the SciPy implementation [56] precomputed at the creation of the archive (L.15);
- (3) looks for the adjacent cells $A[\theta_c].adj$ (using a Delaunay triangulation [14] with the SciPy implementation [56] precomputed at the creation of the archive (L.14)) and extract their task parameters θ and solutions x (L.37-38);
- (4) performs a linear least squares, $M = (\theta^T \theta)^{-1} \theta^T x$ (L.39);
- (5) computes the candidate solution with additional noise (L.40), $x = M \cdot \theta + \sigma_{reg} \cdot \mathcal{N}(0, \text{variance}(x))$ (where σ_{reg} regulates the intensity of the Gaussian noise.)

We use adjacency (i.e., Delaunay triangulation [14]) instead of distance (e.g., the k closest neighbors or the neighbors closer than ϵ) because it adapts the number of adjacent cells to the dimension, thus removing the need to tune a hyperparameter (i.e., the number of neighbors k or the distance ϵ). The regression’s noise coefficient σ_{reg} balances exploration and exploitation. During preliminary experiments, we found that different problems require different optimal values. We concluded on a standard value of 1, even though fine-tuning the value for each problem or having an adaptive selection method could benefit the performances.

4.4 PT-ME Algorithm

Algorithm 1 details PT-ME’s pseudo-code¹. The two major changes with MT-ME [45] are highlighted in colors and summarized to (1) a uniform sampling at each iteration of a new task parameter on which to evaluate a candidate solution and (2) a new variation operator using local linear regression used for 50% of the iterations.

4.5 Distillation

PT-ME asymptotically covers the whole task parameter space by visiting a new task at each iteration. To solve the parametric-task

Algorithm 1 Parametric-Task MAP-Elites

Contributions: Parametric-Task New variation operator

```

1: Parameters:
2:  $d_x$ : dimension of the solution space  $\mathcal{X} = [0, 1]^{d_x}$ 
3:  $d_\theta$ : dimension of the task parameter space  $\Theta = [0, 1]^{d_\theta}$ 
4:  $\text{fitness} : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$ 
5: Hyper-parameters:
6:  $B$ : budget of evaluations (e.g., 100 000)
7:  $N$ : number of cells (e.g., 200)
8:  $S$ : possible tournament sizes (e.g., [1, 5, 10, 50, 100, 500])
9:  $\sigma_{SBX}$ : mutation factor for SBX [1] (e.g., 10.0)
10:  $\sigma_{reg}$ : mutation factor for our new variation operator (e.g., 1.0)
11: Initialization:
12:  $E \leftarrow \emptyset$  ▷  $E$ : to store all the evaluations
13:  $C \leftarrow N$  random centroids using CVT ▷ Divide  $\Theta$  into cells
14:  $D \leftarrow \text{DelaunayTriangulation}(C)$  ▷ For computing adjacency
15:  $kdtree \leftarrow \text{KDTree}(C)$  ▷ To quickly find the cell of a task
16:  $A \leftarrow \emptyset$  ▷  $A$ : Archive
17: for  $\theta_c$  in  $C$  do ▷ one for each cell
18:    $A[\theta_c].\theta = \theta_c$  ▷  $\theta$ : the task parameter
19:    $A[\theta_c].x = \text{random.uniform}(0, 1, )$  ▷  $x$ : the solution
20:    $A[\theta_c].f = \text{fitness}(A[\theta_c].x, A[\theta_c].\theta)$  ▷  $f$ : the fitness
21:    $A[\theta_c].adj = D[\theta_c]$  ▷  $adj$ : the adjacent cells
22: Main loop:
23:  $s \leftarrow \text{random\_in\_list}(S)$  ▷  $s$ : tournament size
24:  $selected \leftarrow \text{zeros}(\text{len}(S))$  ▷ Counter of selection for each size
25:  $successes \leftarrow \text{zeros}(\text{len}(S))$  ▷ Counter of successes for each size
26: for  $_ = 0$  to  $B - N$  do ▷ We already consumed  $N$  evaluations
27:   if  $\text{random}() \leq 0.5$  then ▷ Use SBX with tournament (50%)
28:      $p1, p2 \leftarrow$  two elites randomly selected from  $A$ 
29:      $\theta_{1:s} \leftarrow \text{random.uniform}(0, 1, (s, d_\theta))$  ▷ Candidate tasks
30:      $selected[s] \leftarrow selected[s] + 1$  ▷ Update the bandit
31:      $\theta \leftarrow \text{closest}(\theta_{1:s}, p1.\theta)$  ▷ Task parameters tournament
32:      $\theta_c \leftarrow kdtree.\text{find}(\theta)$  ▷ Find the corresponding cell
33:      $x \leftarrow \text{SBX}(p1, p2, \sigma_{SBX})$  ▷ Mutation & Cross-over
34:   else ▷ Use new variation operator: Local Linear Regression (50%)
35:      $\theta \leftarrow \text{random.uniform}(0, 1, d_\theta)$ 
36:      $\theta_c \leftarrow kdtree.\text{find}(\theta)$  ▷ Find the corresponding cell
37:      $\theta \leftarrow \{A[\theta_i].\theta\}_{\theta_i \in A[\theta_c].adj}$ 
38:      $x \leftarrow \{A[\theta_i].x\}_{\theta_i \in A[\theta_c].adj}$ 
39:      $M \leftarrow (\theta^T \theta)^{-1} \theta^T x$  ▷ Linear least squares
40:      $x \leftarrow M \cdot \theta + \sigma_{reg} \cdot \mathcal{N}(0, \text{variance}(x))$  ▷ Regression + noise
41:      $f \leftarrow \text{fitness}(x, \theta)$  ▷ Evaluate
42:      $E.\text{append}((\theta, x, f))$  ▷ Store the evaluation
43:     if  $f \geq A[\theta_c].f$  then ▷ Update the archive
44:        $A[\theta_c].\theta = \theta$ 
45:        $A[\theta_c].x = x$ 
46:        $A[\theta_c].f = f$ 
47:     if tournament was used then ▷ Update the bandit
48:        $successes[s] \leftarrow successes[s] + 1$ 
49:     if tournament was used then ▷ Update the tournament size
50:        $s \leftarrow S \left[ \underset{j}{\text{argmax}} \left( \frac{successes[j]}{selected[j]} + \sqrt{\frac{2 \ln(\text{sum}(selected))}{selected[j]}} \right) \right]$ 
51: return  $E$ 

```

optimization problem with a finite budget of evaluations and return a solution for each task parameter, we propose to learn an approximation \hat{G} of the function $G : \theta \mapsto x_\theta^*$. To do so, Alg. 2 first computes an archive A_n of elites for a different resolution n using

¹Source code: <https://zenodo.org/doi/10.5281/zenodo.10926438>

Algorithm 2 Compute an archive with a new resolution

```

1: Parameters:
2:  $d_\theta$ : dimension of the task parameter space  $\Theta = [0, 1]^{d_\theta}$ 
3:  $E$ : list of evaluations  $(\theta, x, f)$ 
4:  $N$ : number of cells in the new archive
5: Initialization:
6:  $C \leftarrow N$  random centroids using CVT
7:  $kdtree \leftarrow$  KDTree( $C$ )
8:  $A \leftarrow \{\theta_c : \{\theta : None, x : None, f : 0\} \text{ for } \theta_c \text{ in } C\}$ 
9: Main loop:
10: for  $(\theta, x, f)$  in  $E$  do ▷ Compute elites
11:    $\theta_c \leftarrow kdtree.find(\theta)$  ▷ Find the corresponding cell
12:   if  $f \geq A[\theta_c].f$  then ▷ We suppose  $f$  normalized so that  $f \geq 0$ 
13:      $A[\theta_c] = (\theta, x, f)$ 
14: return  $A$ 

```

the dataset of evaluations returned by PT-ME. We then distill A_n 's elites into a multi-layer perceptron (with the same architecture as PPO's stable-baseline [51, 53], i.e., two hidden layers of 64 neurons) using a supervised regression training with the MSE loss.

5 EXPERIMENTS

5.1 Considered Domains

5.1.1 10-DoF Arm. This toy problem is presented in MT-ME [45] and reused in [33, 57]. It consists in finding the joint positions for different 10-DoF arm morphologies that put the end effector as close to a fixed target (Fig. 2.a).

The task parameter space Θ corresponds to the maximal angular position α_{\max} and the length L of each joint. The solution space \mathcal{X} corresponds to the angular positions of each joint (normalized by the maximal angular position α_{\max}). The fitness f corresponds to the Euclidean distance between the end effector at position pos_{ee} and the fixed target at position pos_{tar} . To get a maximization problem, we put this distance into a negative exponential: $f = \exp(-||pos_{ee} - pos_{tar}||^2)$ so that it is bounded by 0 and 1.

5.1.2 Archery. This toy problem corresponds to shooting an arrow at a range target and computing the associated score (Fig. 3.a).

The task parameter space Θ corresponds to the distance D to the target (between 5m and 40m), and the strength and orientation of horizontal wind W , simulated by an additional constant acceleration (between $-10\text{m}\cdot\text{s}^{-2}$ and $10\text{m}\cdot\text{s}^{-2}$). The solution space \mathcal{X} corresponds to the *yaw* and *pitch* of the arrow (between $-\frac{\pi}{12}$ rad and $\frac{\pi}{12}$ rad) with a constant velocity v of $70\text{m}\cdot\text{s}^{-1}$. The fitness f corresponds to the traditional archery scoring normalized between 0 and 1. More formally, we compute the velocity $\vec{v} = \begin{pmatrix} -\sin(yaw) \\ \cos(yaw)\cos(pitch) \\ \cos(yaw)\sin(pitch) \end{pmatrix}$,

deduce the time of impact $t = \frac{D}{v_y}$, compute the distance to the target $d = ||\frac{1}{2}(-g\vec{z} + W\vec{x})t^2 + \vec{v}t||^2$ and $f = \max(0, \text{int}(10 - \frac{d}{0.061}))/10$.

5.1.3 Door-Pulling. A simulated humanoid robot opens a door by pulling a vertical handle (Fig. 4.a).

The task parameter space Θ corresponds to the position of the door relative to the robot (between 0.8m and 1.2m in front of the robot and 0m and -0.5m laterally) and its orientation (between $\frac{3\pi}{8}$ rad and $\frac{5\pi}{8}$ rad). The fitness f corresponds to the angle of the

door after pulling. This fitness is sparse as the door angle stays null if the robot does not touch the handle. The solution space \mathcal{X} corresponds to the following top-level command sent to a whole-body controller (WBC):

- the gripper's translation from the handle of the door (3DoF);
- the gripper's orientation with constant pitch (2DoF);
- the gripper's translation during pulling (3DoF);
- the gripper's vertical rotation during pulling (1DoF).

The WBC solves at 1 000 Hz a quadratic programming problem that optimizes the robot joint positions to satisfy the quasi-static dynamics while trying to satisfy the top-level commands [13].

Door-Pulling is more similar to a DRL problem, where a policy would be trained to pull the door open from a random initial state. Instead, we explicitly optimize to find top-level commands that pull the door open for every possible configuration of the door with the goal of generalizing those solutions to the whole space.

5.2 Methodology

5.2.1 Baselines.

- **Random** uniformly samples the task parameter θ and the candidate solution x ("How hard is the problem?");
- **CMA-ES** iterates CMA-ES [30] on a random task parameter till it reaches a stopping criterium, then switches to another till it reaches the budget B ("How good is CMA-ES?");
- **CMA-ES (max 10)** does the same but enforces a maximum of 10 iterations per task parameter before switching to another ("Does CMA-ES gain from seeing more different tasks?");
- **PPO** runs the stable-baseline [51] PPO [53] with a random task parameter as initial state, the solution x as action, and the fitness f as reward ("Can DRL solve the problem?");
- **MT-ME** runs MT-ME [45] with 5 000 fixed tasks sampled with CVT [15] ("Do we need a new algorithm?").

5.2.2 PT-ME variants.

- **PT-ME (ours)** uses Alg. 1;
- **PT-ME (no regression)** uses only SBX with tournament ("Is our new variation operator necessary?");
- **PT-ME (100% regression)** uses only the local linear regression ("Is our new variation operator sufficient?");
- **PT-ME (no regression - no tournament)** uses only SBX with no tournament ("Is the tournament useful?");
- **PT-ME (no tournament)** uses both variation operators with no tournament ("Is the tournament useful with the addition of our new variation operator?").

5.2.3 Evaluations and Measures. After a budget B of 100 000 evaluations, we want to evaluate both the coverage and the quality of the solutions with the idea that a perfect algorithm should find optimal solutions evenly spread on the task parameter space. We use the QD-Score introduced in [44] for a finite set of tasks:

$$\text{QD-Score}(A) = \sum_{(\theta, x, f)_{\theta_c} \in A} f$$

where $(\theta, x, f)_{\theta_c}$ is the elite in the cell of centroid θ_c .

As PT-ME stores all the evaluations E from B different tasks (L.42), we recompute the elites for different resolutions, ranging from one cell (the best-found solution on all tasks) up to B cells evenly spread on the task parameter space Θ (we cannot fill more cells than the




Figure 2: (a) Schematic view of the 10-DoF Arm problem. (b) Archive examples extracted from one run. (c) QD-Score for different resolutions (line = median of 20 replications and shaded area = first and third quartiles). (d) Multi-Resolution QD-Score.

budget). We attribute the lowest fitness to each empty cell, meaning that looking at the QD-Score for the highest resolution does not give a complete idea of the algorithm’s performance. To simplify the comparison, we introduce the Multi-Resolution QD-Score (MR-QD-Score) as the mean QD-Score for archives recomputed using Alg. 2 with a range of resolutions from 1 to B cells. More formally:

$$\text{MR-QD-Score}(E, N_{1:n}) = \frac{1}{n} \sum_{i=1}^n \text{QD-Score}(\text{Alg. 2}(E, N_i)) \quad (2)$$

where $N_{1:n}$ are different resolutions. We used $n = 50$ resolutions evenly spread in logspace between 1 and 100 000 cells.

To ease the comparison between problems, we set the fitness minimal value to 0 and its maximal value to 1. It is already the case for Archery (i.e., the arrow can reach the center of the target, $f = 1$, or miss the target, $f = 0$). For the 10-DoF Arm, the fitness is bounded by 0 and 1. However, each task has different bounds (e.g., the arm can be too short to reach the target, leading to a maximal fitness strictly inferior to 1). Knowing the bounds of the fitness functions

for each task parameter is not trivial. We ran CMA-ES [30] on each cell for the highest resolution to estimate them.

All statistical tests are performed with a Mann–Whitney U test and the Bonferroni correction.

5.3 Solutions Evaluations

5.3.1 10-DoF Arm. Fig. 2 shows the results in the 10-DoF Arm problem of all methods with 20 replications each.

CMA-ES allocates too many evaluations for the same task, which prevents it from covering the space and leads to a lower QD-Score than **Random** when there are more than 20 cells to fill. Bounding the number of iterations to 10 allows **CMA-ES (10 max)** to fill more cells with a better solution than **Random**, but it worsens after 1 000 cells. **MT-ME** finds solutions with quality similar to the best methods and outperforms **CMA-ES** and **CMA-ES (10 max)**, showing that solving all the tasks together is more efficient than solving them separately. However, due to its fixed set of tasks, it cannot fill more cells in higher resolutions, significantly decreasing




Figure 3: (a) Schematic view of the Archery problem. (b) Archive examples extracted from one run. (c) QD-Score for different resolutions (line = median of 20 replications and shaded area = first and third quantiles). (d) Multi-Resolution QD-Score.

its QD-Score after 5 000 cells. PPO and all variants of PT-ME have a similar QD-Score profile, succeeding in filling most of the cells with reasonable solutions, even in high resolutions. PT-ME and PPO have the highest MR-QD-Score and are not significantly different.

PT-ME’s MR-QD-Score is significantly higher (p-value<0.001) than all the ablations, showing that the algorithm needs the local linear regression, SBX with a bandit-adapted tournament, and the parametric-task reformulation. Both PT-ME (no tournament) and PT-ME (no regression - no tournament) are significantly worse (p-value<0.001) than PT-ME and PT-ME (no regression), showing that the tournament is beneficial.

Fig. 2.b shows archives for the different methods. CMA-ES and MT-ME have sparser archives as they allocate several evaluations per task parameter. Random fills all task parameter space but does not find solutions as good as PPO and PT-ME and its ablations.

5.3.2 Archery. We then compare all methods on the Archery problem (Fig. 3.c) with 20 replications for each. Compared to 10-DoF Arm, where each candidate solution achieves a positive fitness, for

Archery, most (i.e., approximately 95%) of the solution space leads to a fitness of 0, leading to Random having the worst performance.

CMA-ES and CMA-ES (10 max) have similar QD-Score profiles because either the first generation of candidate solutions has a null fitness and the algorithm stops, or one of them reaches the target by chance, and CMA-ES quickly finds the optimal solution.

MT-ME outperforms CMA-ES and CMA-ES (10 max), showing that solving all the tasks together is more efficient than solving them separately. For resolutions inferior to 10 000 cells, MT-ME is better than PPO, which fails in this problem due to its poor exploration with a sparse reward.

PT-ME is significantly better (p-value<0.001) than all the other methods except PT-ME (100% regression) whose variance is one order of magnitude larger than all the other methods. As it is a pure exploitation method relying on “luck” to explore all space, it can easily get stuck with solutions that never expand the solved region.

Fig. 3.b shows archives extracted from the executions of the different methods. CMA-ES, CMA-ES (10 max), and PPO only solve the easiest tasks (i.e., where the target is the closest). MT-ME

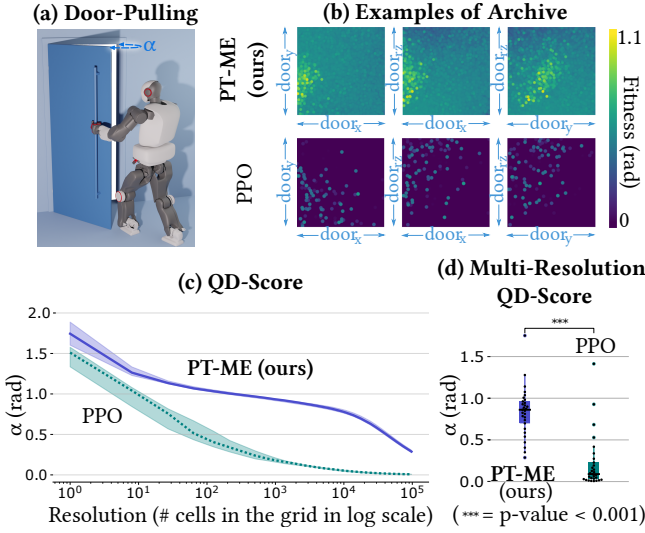


Figure 4: (a) Door-Pulling visualization. (b) Examples of the archive found by PT-ME and PPO [53]. (c) The QD-Score of the two methods for different archive resolutions (the line is the median of 10 replications, and the shaded area is between the first and third quantiles). (d) The Multi-Resolution QD-Score of the two methods for 10 replications.

and PT-ME’s ablations have a similar profile with worse solutions where the wind is the strongest and the target is the farthest.

5.3.3 Door-Pulling. Due to the computing cost of running the humanoid simulation (in the order of the dozens of seconds compared to the order of the millisecond for 10-DoF Arm and Archery), we only compared the result of **PT-ME** and **PPO** [53] on ten replications. The result (Fig. 4) shows that our method significantly outperforms **PPO** (p -value <0.001), which can be explained again by the sparse fitness and **PPO**’s poor exploration abilities.

5.4 Distillation Evaluations

We compare **MT-ME** (for one resolution) and **PT-ME** (for different resolutions) with distillation and **PPO** inference capacities by evaluating the fitness of their proposed solution for 10 000 (1 000 for Door-Pulling) new tasks sampled with CVT [15] (Fig. 5). **PT-ME** significantly outperforms **PPO** on all the three problems (p -value <0.001).

PT-ME inference score first increases as the resolution of the archive increases, showing the need for a large dataset for distillation. The decrease in inference score for large resolutions (i.e., $\geq 10\,000$ cells) is due to the elites’ decreases in quality as the selection pressure decreases. Running the algorithm for longer would increase the maximal resolutions with high-quality elites, even though for Archery, **PT-ME** reaches nearly a perfect score (the lowest run had an inference score of 0.997). For the two toy problems, **PT-ME** outperforms **MT-ME** for resolutions superior to 500 cells and reaches nearly perfect scores, showing that it covers the space with better solutions and solves the parametric-task optimization problem. For Door-Pulling, the inference score is significantly lower

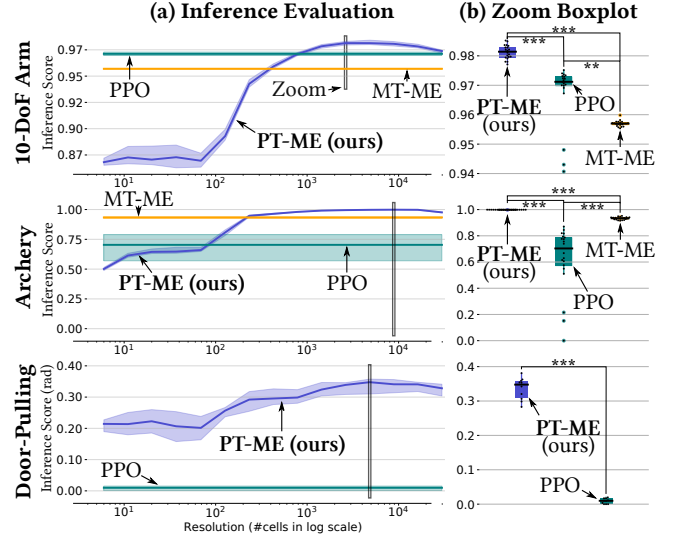


Figure 5: (a) Inference score, i.e., mean fitness over 10 000 (1 000 for Door-Pulling) evenly spread new tasks of PT-ME distillation, MT-ME distillation, and PPO for the three parametric-task optimization problems (the line is the median of 20 (10 for Door-Pulling) replications, and the shaded area is between the first and third quantiles). (b) Box plots with PT-ME’s best resolution (***= p -value <0.001 , **= p -value <0.01).

than the MR-QD-Score due to the chaotic behavior of the contact dynamics of our simulator.

6 DISCUSSION AND CONCLUSION

PT-ME finds high-performing solutions for as many tasks as its budget allows it, making it easy to generalize and distill those solutions into a function approximation that effectively proposes a solution for each task parameter. Its inference significantly outperforms the DRL method PPO [53] on all three problems.

Can PT-ME be used for large task parameter spaces? PT-ME discretizes the task space using CVT [15], which scales up to large dimensions [43]. However, the archive needs more cells to keep a high resolution in large dimensions, which can slow down the convergence. Our new variation operator uses a local linear regression, which could be affected by a larger task space (i.e., the linear model requires the inversion of a matrix). Future work could study other supervised methods (e.g., neural networks) to perform regression.

Can PT-ME be used for large solution spaces? Future work could study the integration of two variation operators to leverage gradient information for scaling up MAP-Elites algorithms to large solution spaces (e.g., a policy neural network): leveraging differentiable problems by computing a gradient from the fitness function [20] (which requires differentiable problems) and training a differentiable critic and exploits it [46] (applicable to black-box problems). Such gradient-based variation operators could be incorporated into our parametric-task framework.

ACKNOWLEDGMENTS

This project is supported by the CPER SCARAT, the CPER CyberEntreprises, the Direction General de l'Armement (convention Inria-DGA "humanoïde résilient"), the Creativ'Lab platform of Inria/LORIA, the EurROBIN Horizon project (grant number 101070596), the Agence de l'Innovation de Défense (AID), and the ANR in the France 2030 program through project PEPR O2R AS3 (ANR-22-EXOD-007).

REFERENCES

- [1] Ram Agrawal, Kalyanmoy Deb, and Ram Agrawal. 2000. Simulated Binary Crossover for Continuous Search Space. *Complex Systems* 9 (06 2000).
- [2] Timothée Anne and Jean-Baptiste Mouret. 2023. Multi-Task Multi-Behavior MAP-Elites. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation* (Lisbon, Portugal) (GECCO '23 Companion). Association for Computing Machinery, New York, NY, USA, 111–114. <https://doi.org/10.1145/3583133.3590730>
- [3] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. 2017. Deep Reinforcement Learning: A Brief Survey. *IEEE Signal Processing Magazine* 34, 6 (2017), 26–38. <https://doi.org/10.1109/MSP.2017.2743240>
- [4] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning* 47 (05 2002), 235–256. <https://doi.org/10.1023/A:1013689704352>
- [5] Kavitesh Kumar Bali, Abhishek Gupta, Yew-Soon Ong, and Puay Siew Tan. 2021. Cognizant Multitasking in Multiobjective Multifactorial Evolution: MO-MFEA-II. *IEEE Transactions on Cybernetics* 51, 4 (2021), 1784–1796. <https://doi.org/10.1109/TCYB.2020.2981733>
- [6] Kavitesh Kumar Bali, Yew-Soon Ong, Abhishek Gupta, and Puay Siew Tan. 2020. Multifactorial Evolutionary Algorithm With Online Transfer Parameter Estimation: MFEA-II. *IEEE Transactions on Evolutionary Computation* 24, 1 (2020), 69–83. <https://doi.org/10.1109/TEVC.2019.2906927>
- [7] Adrien Baranes and Pierre-Yves Oudeyer. 2013. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems* 61, 1 (2013), 49–73.
- [8] S. Barnett. 1968. A simple class of parametric linear programming problems. *Operations Research* 16, 6 (1968), 1160–1165.
- [9] Eric Brochu, Vlad M. Cora, and Nando de Freitas. 2010. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. *CoRR* abs/1012.2599 (2010). arXiv:1012.2599 <http://arxiv.org/abs/1012.2599>
- [10] Wenxue Chen, Changsheng Gao, and Wuxing Jing. 2023. Proximal policy optimization guidance algorithm for intercepting near-space maneuvering targets. *Aerospace Science and Technology* 132 (2023), 108031.
- [11] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (2015), 503–507. <https://doi.org/10.1038/nature14422>
- [12] Weijing Dai, Zhenkun Wang, and Ke Xue. 2022. System-in-package design using multi-task memetic learning and optimization. *Memetic Computing* 14 (03 2022), 1–15. <https://doi.org/10.1007/s12293-021-00346-5>
- [13] Eloïse Dalin, Ivan Bergonzani, Timothée Anne, Serena Ivaldi, and Jean-Baptiste Mouret. 2021. Whole-body teleoperation of the Talos humanoid robot: preliminary results. In *ICRA 2021 - 5th Workshop on Teleoperation of Dynamic Legged Robots in Real Scenarios*. Xi'an / Virtual, China. <https://hal.inria.fr/hal-03245005>
- [14] Boris Delaunay. 1928. Sur la sphere vide. In *Proceedings of the Mathematics, Toronto*. Toronto, 695–700. 11-16 August 1924.
- [15] Qiang Du, Vance Faber, and Max Gunzburger. 1999. Centroidal Voronoi Tessellations: Applications and Algorithms. *SIAM Rev.* 41, 4 (1999), 637–676. <https://doi.org/10.1137/S0036144599352836> arXiv:<https://doi.org/10.1137/S0036144599352836>
- [16] Pinky Dua, K Kouramas, Vivek Dua, and Efstratios N Pistikopoulos. 2008. MPC on a chip—Recent advances on the application of multi-parametric model-based control. *Computers and Chemical Engineering* 32, 4 (2008), 754–765. <https://doi.org/10.1016/j.compchemeng.2007.03.008> Festschrift devoted to Rex Reklaitis on his 65th Birthday.
- [17] Liang Feng, Yuxiao Huang, Lei Zhou, Jinghui Zhong, Abhishek Gupta, Ke Tang, and Kay Chen Tan. 2021. Explicit Evolutionary Multitasking for Combinatorial Optimization: A Case Study on Capacitated Vehicle Routing Problem. *IEEE Transactions on Cybernetics* 51, 6 (2021), 3143–3156. <https://doi.org/10.1109/TCYB.2019.2962865>
- [18] Anthony V. Fiacco. 1976. Sensitivity Analysis for Nonlinear Programming Using Penalty Methods. *Math. Program.* 10, 1 (dec 1976), 287–311. <https://doi.org/10.1007/BF01580677>
- [19] Roger Fletcher. 1987. *Practical Methods of Optimization* (second ed.). John Wiley & Sons, New York, NY, USA.
- [20] Matthew C. Fontaine and Stefanos Nikolaidis. 2021. Differentiable Quality Diversity. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 10040–10052. https://proceedings.neurips.cc/paper_files/paper/2021/file/532923f11ac97d3e7cb0130315b067dc-Paper.pdf
- [21] Matthew C. Fontaine and Stefanos Nikolaidis. 2023. Covariance Matrix Adaptation MAP-Annealing. In *Proceedings of the Genetic and Evolutionary Computation Conference* (Lisbon, Portugal) (GECCO '23). Association for Computing Machinery, New York, NY, USA, 456–465. <https://doi.org/10.1145/3583131.3590389>
- [22] Matthew C. Fontaine, Julian Togelius, Stefanos Nikolaidis, and Amy K. Hoover. 2020. Covariance Matrix Adaptation for the Rapid Illumination of Behavior Space. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference* (Cancún, Mexico) (GECCO '20). Association for Computing Machinery, New York, NY, USA, 94–102. <https://doi.org/10.1145/3377930.3390232>
- [23] Vincent François-Lavet, Peter Henderson, Riashad Islam, Marc G. Bellemare, and Joelle Pineau. 2018. An Introduction to Deep Reinforcement Learning. *Foundations and Trends® in Machine Learning* 11, 3-4 (2018), 219–354. <https://doi.org/10.1561/22000000071>
- [24] Tomas Gal and Josef Nedoma. 1972. Multiparametric linear programming. *Management Science* 18, 7 (1972), 406–422.
- [25] Daniele Gravina, Ahmed Khalifa, Antonios Liapis, Julian Togelius, and Georgios N. Yannakakis. 2019. Procedural Content Generation through Quality Diversity. *2019 IEEE Conference on Games (CoG)* (2019), 1–8. <https://api.semanticscholar.org/CorpusID:195848208>
- [26] Yang Guan, Yangang Ren, Shengbo Eben Li, Qi Sun, Laiquan Luo, and Keqiang Li. 2020. Centralized Cooperation for Connected and Automated Vehicles at Intersections by Proximal Policy Optimization. *IEEE Transactions on Vehicular Technology* 69, 11 (2020), 12597–12608. <https://doi.org/10.1109/TVT.2020.3026111>
- [27] Abhishek Gupta and Yew-Soon Ong. 2018. Memetic Computation: The Mainspring of Knowledge Transfer in a Data-Driven Optimization Era.
- [28] Abhishek Gupta, Yew-Soon Ong, and Liang Feng. 2016. Multifactorial Evolution: Toward Evolutionary Multitasking. *IEEE Transactions on Evolutionary Computation* 20, 3 (June 2016), 343–357. <https://doi.org/10.1109/TEVC.2015.2458037> Conference Name: IEEE Transactions on Evolutionary Computation.
- [29] Abhishek Gupta, Lei Zhou, Yew-Soon Ong, Zefeng Chen, and Yaqing Hou. 2022. Half a Dozen Real-World Applications of Evolutionary Multitasking, and More. *IEEE Computational Intelligence Magazine* 17 (05 2022), 49–66. <https://doi.org/10.1109/MCI.2022.3155332>
- [30] Nikolaus Hansen, Sibylle Müller, and Petros Koumoutsakos. 2003. Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES). *Evolutionary computation* 11 (02 2003), 1–18. <https://doi.org/10.1162/106365603321828970>
- [31] Xingxing Hao, Rong Qu, and Jing Liu. 2021. A Unified Framework of Graph-Based Evolutionary Multitasking Hyper-Heuristic. *IEEE Transactions on Evolutionary Computation* 25, 1 (2021), 35–47. <https://doi.org/10.1109/TEVC.2020.2991717>
- [32] Shijia Huang, Jinghui Zhong, and Wei-Jie Yu. 2021. Surrogate-Assisted Evolutionary Framework with Adaptive Knowledge Transfer for Multi-Task Optimization. *IEEE Transactions on Emerging Topics in Computing* 9, 4 (2021), 1930–1944. <https://doi.org/10.1109/TETC.2019.2945775>
- [33] Binh Huynh Thi Thanh, Le Van Cuong, Ta Bao Thang, and Nguyen Hoang Long. 2023. Ensemble Multifactorial Evolution With Biased Skill-Factor Inheritance for Many-Task Optimization. *IEEE Transactions on Evolutionary Computation* 27, 6 (2023), 1735–1749. <https://doi.org/10.1109/TEVC.2022.3227120>
- [34] Jeppe Theiss Kristensen and Paolo Burelli. 2020. Strategies for Using Proximal Policy Optimization in Mobile Puzzle Games. In *Proceedings of the 15th International Conference on the Foundations of Digital Games* (Bugibba, Malta) (FDG '20). Association for Computing Machinery, New York, NY, USA, Article 2, 10 pages. <https://doi.org/10.1145/3402942.3402944>
- [35] Joel Lehman and Kenneth Stanley. 2011. Abandoning Objectives: Evolution Through the Search for Novelty Alone. *Evolutionary computation* 19 (06 2011), 189–223. https://doi.org/10.1162/EVCO_a_00025
- [36] Jing Liang, Kangjia Qiao, Minghua Yuan, Kunjie Yu, Boyang Qu, Shilei Ge, Yaxin Li, and Guanlin Chen. 2020. Evolutionary multi-task optimization for parameters extraction of photovoltaic models. *Energy Conversion and Management* 207 (03 2020), 112509. <https://doi.org/10.1016/j.enconman.2020.112509>
- [37] Zhengping Liang, Xiujun Xu, Ling Liu, Yaofeng Tu, and Zexuan Zhu. 2022. Evolutionary Many-Task Optimization Based on Multisource Knowledge Transfer. *IEEE Transactions on Evolutionary Computation* 26, 2 (2022), 319–333. <https://doi.org/10.1109/TEVC.2021.3101697>
- [38] Rung-Tzuo Liaw and Chuan-Kang Ting. 2017. Evolutionary many-tasking based on biocoenosis through symbiosis: A framework and benchmark problems. In *2017 IEEE Congress on Evolutionary Computation (CEC)*. 2266–2273. <https://doi.org/10.1109/CEC.2017.7969579>
- [39] Siyu Lin and Peter A. Beling. 2021. An end-to-end optimal trade execution framework based on proximal policy optimization. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*. 4548–4554.

- [40] Junwei Liu, Peiling Li, Guibin Wang, Yongxing Zha, Jianchun Peng, and Gang Xu. 2020. A Multitasking Electric Power Dispatch Approach With Multi-Objective Multifactorial Optimization Algorithm. *IEEE Access* 8 (2020), 155902–155911. <https://doi.org/10.1109/ACCESS.2020.3018484>
- [41] Songrit Maneewongvatana and David M Mount. 1999. Analysis of approximate nearest neighbor searching with clustered point sets. *arXiv preprint cs/9901013* (1999).
- [42] Alan Tan Wei Min, Yew-Soon Ong, Abhishek Gupta, and Chi-Keong Goh. 2019. Multiproblem Surrogates: Transfer Evolutionary Multiobjective Optimization of Computationally Expensive Problems. *IEEE Transactions on Evolutionary Computation* 23, 1 (2019), 15–28. <https://doi.org/10.1109/TEVC.2017.2783441>
- [43] Jean-Baptiste Mouret. 2023. Fast generation of centroids for MAP-Elites. In *Companion Proceedings of the Conference on Genetic and Evolutionary Computation, GECCO 2023, Companion Volume, Lisbon, Portugal, July 15–19, 2023*, Sara Silva and Luis Paquete (Eds.). ACM, 155–158. <https://doi.org/10.1145/3583133.3590726>
- [44] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. <https://doi.org/10.48550/ARXIV.1504.04909>
- [45] Jean-Baptiste Mouret and Glenn Maguire. 2020. Quality Diversity for Multi-Task Optimization. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference (Cancún, Mexico) (GECCO '20)*. Association for Computing Machinery, New York, NY, USA, 121–129. <https://doi.org/10.1145/3377930.3390203>
- [46] Olle Nilsson and Antoine Cully. 2021. Policy gradient assisted MAP-Elites. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Lille France, 866–875. <https://doi.org/10.1145/3449639.3459304>
- [47] Michael Pearce and Juergen Branke. 2018. Continuous multi-task Bayesian Optimisation with correlation. *European Journal of Operational Research* 270, 3 (2018), 1074–1085. <https://doi.org/10.1016/j.ejor.2018.03.017>
- [48] Luca Pinciroli, Piero Baraldi, Guido Ballabio, Michele Compare, and Enrico Zio. 2021. Deep Reinforcement Learning Based on Proximal Policy Optimization for the Maintenance of a Wind Farm with Multiple Crews. *Energies* 14, 20 (2021). <https://doi.org/10.3390/en14206743>
- [49] Efstratios N. Pistikopoulos, Vivek Dua, Nikolaos A. Bozinis, Alberto Bemporad, and Manfred Morari. 2000. On-line optimization via off-line parametric optimization tools. *Computers & Chemical Engineering* 24, 2 (July 2000), 183–188. [https://doi.org/10.1016/S0098-1354\(00\)00510-X](https://doi.org/10.1016/S0098-1354(00)00510-X)
- [50] Efstratios N Pistikopoulos, Vivek Dua, Nikolaos A Bozinis, Alberto Bemporad, and Manfred Morari. 2002. On-line optimization via off-line parametric optimization tools. *Computers & Chemical Engineering* 26, 2 (2002), 175–185.
- [51] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. 2021. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research* 22, 268 (2021), 1–8. <http://jmlr.org/papers/v22/20-1364.html>
- [52] Ramon Sagarna and Yew-Soon Ong. 2016. Concurrently searching branches in software tests generation through multitask evolution. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*. 1–8. <https://doi.org/10.1109/SSCI.2016.7850040>
- [53] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR abs/1707.06347* (2017).
- [54] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. 2016. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proc. IEEE* 104, 1 (2016), 148–175. <https://doi.org/10.1109/JPROC.2015.2494218>
- [55] Haoran Sun, Linhan Yang, Yuping Gu, Jia Pan, Fang Wan, and Chaoyang Song. 2023. Bridging Locomotion and Manipulation Using Reconfigurable Robotic Limbs via Reinforcement Learning. *Biomimetics* 8, 4 (2023). <https://doi.org/10.3390/biomimetics8040364>
- [56] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, Ilhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- [57] Chao Wang, Jing Liu, Kai Wu, and Zhaoyang Wu. 2022. Solving Multitask Optimization Problems With Adaptive Knowledge Transfer via Anomaly Detection. *IEEE Transactions on Evolutionary Computation* 26, 2 (2022), 304–318. <https://doi.org/10.1109/TEVC.2021.3068157>
- [58] Jian Yin, Anmin Zhu, Zexuan Zhu, Yanan Yu, and Xiaoliang Ma. 2019. Multifactorial Evolutionary Algorithm Enhanced with Cross-task Search Direction. In *2019 IEEE Congress on Evolutionary Computation (CEC)*. 2244–2251. <https://doi.org/10.1109/CEC.2019.8789959>
- [59] Gen Yokoya, Heng Xiao, and Toshiharu Hatanaka. 2019. Multifactorial optimization using Artificial Bee Colony and its application to Car Structure Design Optimization. In *2019 IEEE Congress on Evolutionary Computation (CEC)*. 3404–3409. <https://doi.org/10.1109/CEC.2019.8789940>
- [60] Ming Zhang, Yang Lu, Youxi Hu, Nasser Amaitik, and Yuchun Xu. 2022. Dynamic Scheduling Method for Job-Shop Manufacturing Systems by Deep Reinforcement Learning with Proximal Policy Optimization. *Sustainability* 14, 9 (2022). <https://doi.org/10.3390/su14095177>
- [61] Hong Zhao, Xuhui Ning, Xiaotao Liu, Chao Wang, and Jing Liu. 2023. What makes evolutionary multi-task optimization better: A comprehensive survey. *Applied Soft Computing* 145 (2023), 110545. <https://doi.org/10.1016/j.asoc.2023.110545>
- [62] Jinghui Zhong, Liang Feng, Wentong Cai, and Yew-Soon Ong. 2020. Multifactorial Genetic Programming for Symbolic Regression Problems. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 50, 11 (2020), 4492–4505. <https://doi.org/10.1109/TSMC.2018.2853719>
- [63] Jacques Zhong, Vincent Weistroffer, Jean-Baptiste Mouret, Francis Colas, and Pauline Maurice. 2023. Workstation Suitability Maps: Generating Ergonomic Behaviors on a Population of Virtual Humans With Multi-Task Optimization. *IEEE Robotics Autom. Lett.* 8, 11 (2023), 7384–7391. <https://doi.org/10.1109/LRA.2023.3318191>