

Simultaneously Evolving Deep Reinforcement Learning Models using Multifactorial Optimization

Aritz D. Martinez^{†*}, Eneko Osaba^{†*}, Javier Del Ser^{†,‡} and Francisco Herrera[§]

[†]TECNALIA, Basque Research and Technology Alliance (BRTA), 48160 Derio, Bizkaia, Spain

Email: {aritz.martinez, eneko.osaba, javier.delsers}@tecnalia.com

[‡]University of the Basque Country, 48013 Bilbao, Bizkaia, Spain

[§]DaSCI Andalusian Institute of Data Science and Computational Intelligence. University of Granada. 18071 Granada, Spain

* Corresponding authors. These authors have equally contributed to the work presented in this paper.

Abstract— In recent years, Multifactorial Optimization (MFO) has gained a notable momentum in the research community. MFO is known for its inherent capability to efficiently address multiple optimization tasks at the same time, while transferring information among such tasks to improve their convergence speed. On the other hand, the quantum leap made by Deep Q Learning (DQL) in the Machine Learning field has allowed facing Reinforcement Learning (RL) problems of unprecedented complexity. Unfortunately, complex DQL models usually find it difficult to converge to optimal policies due to the lack of exploration or sparse rewards. In order to overcome these drawbacks, pre-trained models are widely harnessed via Transfer Learning, extrapolating knowledge acquired in a source task to the target task. Besides, meta-heuristic optimization has been shown to reduce the lack of exploration of DQL models. This work proposes a MFO framework capable of simultaneously evolving several DQL models towards solving interrelated RL tasks. Specifically, our proposed framework blends together the benefits of meta-heuristic optimization, Transfer Learning and DQL to automate the process of knowledge transfer and policy learning of distributed RL agents. A thorough experimentation is presented and discussed so as to assess the performance of the framework, its comparison to the traditional methodology for Transfer Learning in terms of convergence, speed and policy quality, and the intertask relationships found and exploited over the search process.

Index Terms—Multifactorial Optimization, Deep Reinforcement Learning, Transfer Learning, Evolutionary Algorithm.

I. INTRODUCTION

Since its inception in the literature mainstream, Deep Learning has been typically used for modeling (i.e. generative distribution learning and predictive approaches) and actuation (correspondingly, Reinforcement Learning, RL). While much attention has been traditionally paid to the former, the irruption of multi-layered complex neural architectures to the RL paradigm marked a milestone in regards to the capability of artificial agents learn to optimally interact with a given environment. Among them, Deep Q Learning (DQL) introduced the use of Deep Neural Networks to predict the total reward expected to be gained after a particular action was taken on the environment. By leveraging this increased reward modeling capability along with other improvements (e.g. experience replay), DQL has since then excelled at learning over challenging RL domains, in some of them attaining near-human levels of performance.

Since the evaluation budget to solve a RL task is often assumed to be unlimited, the complexity of the environment does not have much relevance, given that the algorithm guarantees convergence to a sub-optimal solution. In this context, it has been proven that Deep Reinforcement Learning can achieve impressive results with a proper number of evaluations. In recent works such as the `HideAndSeek` game from OpenAI Labs [1], billions of evaluations are needed to make interesting behaviors emerge among the agents. Nevertheless, in such complex environments, or even in contexts where the amount of evaluations is limited, Transfer Learning is instead used to exploit the knowledge captured by models trained beforehand for other related tasks. This transferred knowledge is adopted as the initial state for the agent to be learned in the target scenario, which allows reducing the time needed to grasp and acquire sophisticated behavioral skills.

Although Transfer Learning has a straightforward application among Deep Learning models, its effectiveness is stringently subject to the similarity between origin and destination tasks [2]. Unfortunately, it is not always the case in which two tasks can be claimed to be similar to each other and suitable for Transfer Learning before actually testing their empirical performance. For instance, in the `HideAndSeek` scenario tackled in [1], Transfer Learning is used to make agents perform different complex tasks departing, as starting state of the Deep RL model, knowledge captured for other tasks such as `ObjectCounting`, `LockAndReturn`, `ShelterConstruction` or `SequentialLock`, among others (i.e. pretrained on these tasks). Despite the benefits of Transfer Learning for the RL realm, determining the transferability of knowledge among tasks is, to date, delegated to common sense and intuition held on the scenarios at hand.

The work presented in this paper aims to take a step forward in this direction. Specifically, we propose a framework for simultaneously learning multiple DQL models for RL tasks, leveraging the existing relationships and complementarities among such tasks during the training phase. To this end, the learning process is formulated as an optimization problem defined on a common search space, from which solutions to specific DQL models for every task can be decoded. Efficiently solving this optimization problem is done by means of Transfer Optimization, an emerging paradigm for tackling different

problems by automatically exploiting synergies among them learned over the optimization process [3]. Several categories can be established under this paradigm (e.g. *sequential transfer* [4]). In this work we select *multitasking* to dynamically optimize multiple DQL models, and to capitalize on the relationships between RL scenarios during the search [5], [6]. In particular, we resort to the so-called Multifactorial Evolutionary Algorithm (MFEA, [7]), a specific multitasking algorithm that relies on evolutionary search operators (also known as Evolutionary Multitasking [8]). The intrinsic knowledge transfer capability between candidates featured by MFEA allows for a simultaneous co-evolution of multiple DQL models, achieving faster convergence and better performance levels by virtue of similarities among RL tasks.

A experimental setup is devised to analyze and validate the performance of the proposed framework when undertaking several simple RL scenarios. As discussed from the obtained results, MFEA not only renders an excellent performance when learning all DQL models at the same time in terms of accuracy and convergence speed, but also evinces whether the search algorithm identifies and exploits the inherent relationships between the defined RL scenarios. The results stemming from this research work can buttress the idea that a unique search process can face the main drawbacks of DQL: i) on the one hand, sharing a unique search space avoids the evolution process to get stuck in local optima; and ii) on the other hand, knowledge is shared between different tasks without any information on whether they are related anyhow, hence acting as an *adaptive* transfer learning mechanism.

The rest of the paper is organized as follows: first, Section II introduces key concepts on MFEA, DQL and Transfer Learning that are required for a proper understanding of the present contribution. Next, Section III describes in detail the proposed framework. Experiments and results are presented and discussed in Section IV. Finally, conclusions and prospects towards future research lines are given in Section V.

II. BACKGROUND

As stated above, this section introduces background concepts related to the present work. To begin with, Subsection II-A formally exposes the MFEA algorithm adopted in this work. Then, Subsection II-B provides an introduction to DQL, while Subsection II-C explains how Transfer Learning is usually performed on RL scenarios.

A. Multifactorial Evolutionary Algorithm

MFEA is a recently proposed evolutionary algorithm for addressing multiple optimization problems using bio-cultural schemes of multifactorial inheritance [7]. Briefly explained, MFEA considers an scenario with K optimization problems (*tasks*), defined as $\{T_1, \dots, T_K\} = \{T_k\}_{k=1}^K$. Without loss of generality we assume that these tasks are minimization problems, each defined over a search space \mathcal{X}_k and $D_k = |\mathcal{X}_k|$ (dimensionality). The goal is to find solutions $\{\mathbf{x}_k^*\}_{k=1}^K$ such that $\mathbf{x}_k^* = \arg_{\mathbf{x} \in \mathcal{X}_k} \min F_k(\mathbf{x}) \forall k \in \{1, \dots, K\}$, where $F_k : \mathcal{X}_k \mapsto \mathbb{R}$ is the objective function of task T_k .

Some concepts must be posed before continuing further. Firstly, it is necessary to define a unified search space \mathcal{X}^U for all tasks, wherein solutions to each of such problems can be encoded and decoded. The form taken by this unified search space \mathcal{X}^U depends roughly on the optimization problems to be solved, with several alternatives contributed to date (e.g. Random Keys [9]). The dimensionality of the unified search space is usually¹ set to be the maximum among the dimensions of the tasks under consideration, i.e. $|\mathcal{X}^U| = \max_k D_k$.

In order to be able to compare among candidate solutions in the unified space, it is necessary to establish a global search method. According to evolutionary computation principles, a population of candidates $\{\mathbf{x}^p\}_{p=1}^P$ is defined, wherein each individual \mathbf{x}^p can be decoded to $\mathbf{x}_k^p \in \mathcal{X}_k$ related to task T_k . Once decoded, several concepts are defined to rank and evolve the individuals within the population:

- Factorial cost ($\Psi_k^p \in \mathbb{R}$): The factorial cost of candidate \mathbf{x}^p in task T_k is given by $\Psi_k^p = \lambda \cdot \delta_k^p + F_k(\mathbf{x}_k^p)$, where λ is a penalizing multiplier, δ_k^p represents the total constraint violation of \mathbf{x}^p at task T_k , and $F_k(\mathbf{x}_k^p)$ denotes the objective value once \mathbf{x}^p has been decoded to \mathbf{x}_k^p . If there is not constraint violation, $\Psi_k^p = F_k(\mathbf{x}_k^p)$.
- Factorial rank ($r_k^p \in \{1, \dots, P\}$) is defined as the position or rank of candidate \mathbf{x}^p in task T_k after sorting the population in ascending order of Ψ_k^p .
- Scalar fitness ($\varphi^p \in \mathbb{R}[0, 1]$) is given by the inverse of the factorial rank r_k^p of candidate \mathbf{x}^p in the task in which it performs best, i.e., $\varphi^p = 1 / (\min_{k \in \{1 \dots K\}} r_k^p)$.
- Skill factor ($\tau^p \in \{1, \dots, K\}$) is the task in which \mathbf{x}^p performs best, i.e., $\tau^p = \arg_{k \in \{1, \dots, K\}} \min r_k^p$.

When the scalar fitness φ^p has been computed $\forall \mathbf{x}^p$, they can be compared to each other by assuming that \mathbf{x}^p dominates $\mathbf{x}^{p'}$ if $\varphi^p > \varphi^{p'}$. Based on this relationship of dominance and the encoding/decoding procedure from \mathcal{X}^U to each \mathcal{X}_k and vice versa, an evolutionary optimization algorithm can be defined, comprising crossover, mutation and selection operators operating on the aforementioned population of individuals.

Algorithm 1: Multifactorial Evolutionary Algorithm

- 1 Initialize at random a population $\{\mathbf{x}^p\}_{p=1}^P$
 - 2 Compute factorial costs $\Psi_k^p \forall k, p$
 - 3 Assign a skill factor τ^p to each \mathbf{x}^p
 - 4 **while** *termination criterion not reached* **do**
 - 5 Apply operators to $\{\mathbf{x}^p\}_{p=1}^P$, yielding $\{\mathbf{x}^{p,\Delta}\}_{p=1}^{P\Delta}$
 - 6 Evaluate each $\mathbf{x}^{p,\Delta}$ on its parents' skill factor
 - 7 Merge $\{\mathbf{x}^{p,\Delta}\}_{p=1}^{P\Delta}$ and $\{\mathbf{x}^p\}_{p=1}^P$
 - 8 Update φ^p and τ^p of merged candidates
 - 9 Select P best candidates in terms of
 - 10 **end**
 - 11 Return the best candidate for each of the K tasks
-

This is indeed the approach followed by MFEA, whose search procedure is described in Algorithm 1. Apart from the

¹This depends on how solutions belonging to the unified search space are decoded to solutions of each problem T_k .

concept of unified search space and the selection based on factorial ranks and factorial costs, MFEA differs from standard evolutionary heuristics also in terms of search operators, featuring novel methods (namely, *Assortative Mating* and *Vertical Cultural Transmission*) that favor the exchange of knowledge among tasks. On one hand, *Assortative Mating* (Algorithm 2) imposes that candidates with the same skill factor mate more likely than two candidates excelling at two different tasks.

Algorithm 2: Assortative mating in MFEA

- 1 Define a threshold $\varphi \in \mathbb{R}[0, 1]$
 - 2 Draw a random number $r \sim \text{Uniform}(0, 1)$
 - 3 Select two parents \mathbf{x}^{p1} and \mathbf{x}^{p2} from the population
 - 4 **if** $\tau^{p1} = \tau^{p2}$ **or** $r < \varphi$ **then**
 - 5 Apply crossover and mutation on \mathbf{x}^{p1} and \mathbf{x}^{p2} to generate two offspring $\mathbf{x}^{p1,\Delta}$ and $\mathbf{x}^{p2,\Delta}$
 - 6 **else**
 - 7 Apply mutation on \mathbf{x}^{p1} to yield $\mathbf{x}^{p1,\Delta}$
 - 8 Apply mutation on \mathbf{x}^{p2} to yield $\mathbf{x}^{p2,\Delta}$
 - 9 **end**
-

On the other hand, *Vertical Cultural Transmission* enforces that candidates are not evaluated in all optimization tasks $\{T_1, \dots, T_K\}$, but only on one of their parents' skill factors. This specific aspect makes MFEA competitive in terms of computational complexity. In addition to these two search operators, MFEA implements a local search method before evaluating a candidate, which we will not consider in our study for reasons disclosed later. Finally, once the produced offspring have been evaluated and merged with those from the previous generation, selection based on scalar fitness φ^p is applied, by which the best P individuals in terms of this fitness are retained in the population for the next generation.

B. Deep Reinforcement Learning

As mentioned in the introduction, the proposed framework deals with distributed RL problems undertaken by means of DQL models, whose learning process is formulated as an optimization problem and solved jointly via evolutionary multitasking. DQL has gained a growing popularity during recent years for their superior modeling [10] and multi-agent cooperation capabilities [1]. However, what makes DQL models particularly interesting for the purpose of this study is their suitability to be trained by using metaheuristic algorithms [11], showing competitive results when compared to traditional back-propagation. For the sake of completeness, we next elaborate briefly on the concepts underneath DQL.

In general, the goal of RL is to learn a policy, usually denoted as $\pi_{\theta}(s)$ (with θ denoting the parameters defining the policy itself), to efficiently perform a task depending on a state $s \in \mathcal{S}$. In general terms, a RL approach is described by an agent in charge of taking an action $a_t = \pi_{\theta}(s_t)$ from a set of actions \mathcal{A} , every time a new state $s_t \in \mathcal{S}$ is presented to it. The environment, where the task is defined, returns a reward of performing action a_t in state s_t , and informs the

agent with the next state s_{t+1} of the environment. The search for a policy $\pi_{\theta}(s)$ capable of maximizing this reward and optimally solving the RL task is not trivial, unchaining a flurry of different algorithmic approaches and heuristics for this purpose [12].

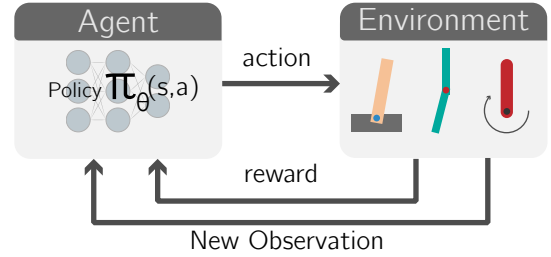


Fig. 1. General architecture of DQL, in which the policy to be learned is implicitly embedded in the trainable parameters of a Deep Neural Network.

One of such different policy learning methods is DQL, in which the policy is represented by the weights and biases of a Deep Neural Network (Figure 1). The Q in DQL comes from the traditional Q learning, and refers to the long-term return $Q(s_t, a_t)$ predicted for every state-action pair at instant t . This value is updated every time the agent takes an action, as per:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [R_t + \gamma \max_{a \in \mathcal{A}} Q(s_{t+1}, a) - Q(s_t, a_t)],$$

where $Q(s_t, a_t)$ denotes the current Q value associated to taking the action a_t at state s_t , R_t is the reward received when moving from state s_t to state s_{t+1} , $\alpha \in \mathbb{R}[0, 1]$ is a learning rate, and γ is a discount factor that lowers the contribution of future rewards in the computation. As opposed to this update rule, DQL resorts to a Deep Neural Network to approximate the Q value associated to every task (outputs) given the current state (input). As a result, the need for computing and updating Q for every (state, action) pair is overridden, making it suitable for RL problems of high dimensionality.

Since the learned policy $\pi_{\theta}(s)$ is implicitly encoded in the weights and biases of the Deep Neural Network (θ), the behavior of the agent when acting on the environment only depends on the values of such neural parameters. To set optimally these values, gradient descent is often utilized² to minimize a Q -dependent squared error loss function with respect to θ :

$$\min_{\theta} \frac{1}{2} \left(M_{\theta}(s_t, a_t) - [R_t + \gamma \max_{a \in \mathcal{A}} M_{\theta}(s_{t+1}, a)] \right)^2, \quad (1)$$

where $M_{\theta}(s, a)$ the output of the neural network with parameters θ corresponding to action a when observing state s .

Indeed, the training problem can be seen as an optimization task wherein parameter values act as decision variables, and the loss function as the objective to be minimized. This fact underlies the reason why metaheuristic algorithms have been proposed to evolve DQL models, with the goal of harnessing

²Due to the lack of space we skip specific aspects of DQL model training such as experience replay, target network and reward clipping.

the highly parallel nature of these solvers towards finding good policies at a lower computational cost [11], [13]–[15]. More recently, alternative techniques such as Novelty Search have been hybridized with metaheuristic methods to reach even better results for evolving DQL models [16]. Neuroevolution for DQL models has been approached with evolutionary algorithms as well [17].

C. Transfer Learning

Finally, a third ingredient involved in our framework is Transfer Learning, which refers to the knowledge transfer between models solving tasks that share some similarities [18]. Conforming to intuition, the most closely tasks are related to each other, the most effective the transfer of knowledge between models will be. This statement also holds in Transfer Optimization, in which the relationship between the problems being solved also drives the potential benefit of exchanging knowledge among them. The main rationale of Transfer Learning when utilized in the context of data modeling tasks is to quickly adapt a model for performing well in different environments, with a lower demand for training data or at a lower training cost.

While Transfer Learning has been a thoroughly addressed matter of study, the simplicity by which Transfer Learning can be realized between Deep Learning models has allowed them to dominate this research arena, particularly in image recognition and reinforcement learning scenarios [19], [20]. The most common approach is to reuse pre-trained models learned for the source task(s) as a initialization point of the model of the target task. To this end, the most straightforward approach is to copy a fraction of the parameter values of the pre-trained model to the model that faces the target task. Once this is done, some layers of the model of the target task are frozen, whereas parameter belonging to the rest of the layers are trained via back-propagation.

Despite this apparent simplicity, it is not trivial to determine how many layers/parameters to transfer between models: the more similar the tasks are, the higher the number of parameters that can be transferred to the model of the target task will be. Furthermore, it still remains unclear whether the degree of similarity between origin and target task can be easily quantified in complex modeling problems, specially in those where no intuition can be held in this matter. In addition, the mapping between the degree of similarity among tasks and the amount of parameter values transferred between models is still open to further research. With all this been said, Transfer Learning can be useful between DQL to develop general controllers that can adapt themselves to different environments, and solve accurately RL problems therein. The framework presented in the next section aims precisely at this goal, further complementing it with its capability to enforce information transfer among RL tasks during the learning process.

III. PROPOSED FRAMEWORK

The framework proposed in this paper simultaneously evolves DQL models aimed at solving distinct albeit inter-

related RL tasks. To perform this efficiently, we adopt MFEA as the method in charge of solving the different task bunch. Besides, a novel encoding scheme is devised, whose design is inspired by the way Transfer Learning is implemented among neural networks.

To begin with, the optimization problems $\{T_1, \dots, T_K\}$ represent RL tasks defined on a set of different environments $\mathcal{E} = \{E_1, \dots, E_N\}$, such that the environment index on which problem T_k is defined is given by $\xi(k)$, with $\xi : \{1, \dots, K\} \mapsto \{1, \dots, N\}$. Furthermore, each of such environments can be *configured* with a set of parameters β that permit to tailor different configurations of the environment for the task to be solved. For instance, in a `cartpole` game, β could represent the gravity force or the length of the pole to be kept upright³. Therefore, tasks to be solved may belong to different environments, wherein tasks defined on the same environment may operate under different configurations. Ideally, knowledge transfer should occur more likely among tasks defined on the same environment (*intra-environment transfer*), as long as the configurations β do not make the tasks radically different to each other. When extrapolating this statement to the knowledge exchange among different environments (corr. *inter-environment transfer*), greater differences among the tasks can be expected, yet still leaving room for knowledge exchange (e.g. levering a pole in `cartpole` should contribute to swinging up a two-link robot in the `acrobot` game).

Automating this knowledge transfer among DQL models aimed to solve different tasks is what motivates the adoption of MFEA. To this end, for each task $k \in \{1, \dots, K\}$ a DQL model $M_{\theta_k}^k$ is sought, whose parameters θ_k (weights, biases) constitute the local decision variables that MFEA aims to optimize via a single exploration over a unified search space. Therefore, according to the notation introduced for MFEA in Subsection II-A, $\theta_k \equiv \mathbf{x}_k$ and \mathcal{X}_k is given by \mathbb{R}^{D_k} , with $D_k = |\theta_k|$. Based on this notation, the problem T_k to be solved is formulated as the maximization of the final reward of the DQL model $M_{\theta_k}^k$ averaged over a number N_{tst} of test episodes run over environment $\xi(k)$:

$$\text{Problem } T_k : \max_{\theta_k \in \mathbb{R}^{D_k}} \frac{1}{N_{tst}} \sum_{t=1}^{N_{tst}} R(M_{\theta_k}^k; E_{\xi(k)}, t), \quad (2)$$

where $R(M; E, t)$ denotes the final reward achieved by model M on environment E at test episode t . We note here that the above objective differs from the one pursued in the traditional DQL training approach, given in Expression (1). The reason is that we do not seek to learn φ_k online over epochs by interacting with the environment, but rather to optimize them off-line. Nevertheless, this clarification stimulates research lines further exposed in Section V.

In order to realize a unified search space \mathcal{X}^U on which to perform the multifactorial search, two important factors must be taken into account when designing the encoding and decoding procedures between \mathcal{X}^U and \mathcal{X}_k : i) they must be

³Further information about these exemplifying games, which are also used in the experiments, can be found in the OpenAI Gym Library [21].

computationally efficient; and ii) they must harness any a priori knowledge on the type of tasks to be solved. Indeed, when Transfer Learning is applied to RL, there are some considerations that can be taken into account in order to obtain a better performance [19]. It is known that neural networks learn general patterns in their first layers, while the last ones retain specialized information about the task under consideration. Consequently, as mentioned in Subsection II-C, Transfer Learning among Deep Neural Networks is mostly made by sharing network parameters corresponding to the first layers of the network, while delegating the specialization for the task to be solved to the parameters of the last layers of the neural architecture.

Bearing these observations in mind, it seems natural to embrace them when designing \mathcal{X}^U . Specifically, we partition any given solution $\mathbf{x}^p \in \mathcal{X}^U$ in two parts:

- A first part \mathcal{X}_{sh}^U representing the parameters of the L_{sh} first layers of all models $\{M_{\theta_k}^k\}_{k=1}^K$, whose values will be the same across tasks; and
- K parts $\{\mathcal{X}_k^U\}_{k=1}^K$, each embedding the parameters of the last layers of $M_{\theta_k}^k$ that are specialized for T_k and hence, not shared across different tasks.

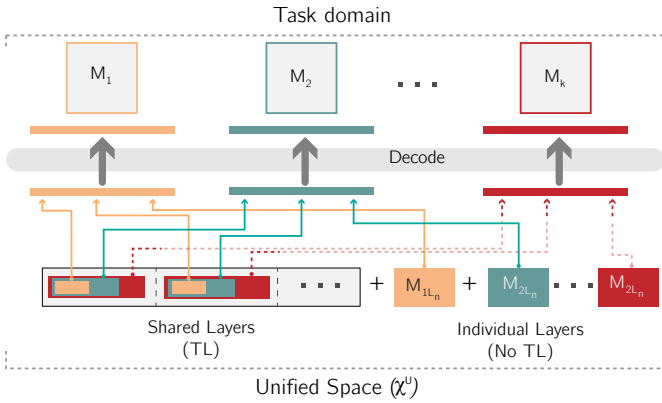


Fig. 2. Encoding and decoding scheme of the proposed MFEA framework.

The decoding process is graphically depicted in Figure 2, where it can be observed that candidate solutions $\theta_k \equiv \mathbf{x}_k$ are built by extracting the shared weights and biases in \mathcal{X}_{sh}^U , followed by those in \mathcal{X}_k^U . Therefore, the dimensionality of the unified search space is given by:

$$|\mathcal{X}^U| = \sum_{l=1}^{L_{sh}} \max_k \{|\varphi_k| : \varphi_k \in \text{layer } l \text{ of } M_{\theta_k}^k\} + \sum_{k=1}^K \sum_{l=L_{sh}+1}^{L^k} |\varphi_k| : \varphi_k \in \text{layer } l \text{ of } M_{\theta_k}^k, \quad (3)$$

where L^k denotes the number of layers of $M_{\theta_k}^k$. Once \mathcal{X}^U has been defined, knowledge among $\{T_k\}_{k=1}^K$ is intrinsically shared via the unified search space and the MFEA evolutionary operators (Algorithm 1). As such, the population of individuals $\{\mathbf{x}^p\}_{p=1}^P$ defined on the unified search space is evolved over

generations by successively applying crossover and mutation operators through the assortative mating procedure of Algorithm 2. This operator favors the implicit genotype transfer between tasks, allowing a RL problems to benefit from the knowledge gained to solve other RL problems embedded in the weights and biases belonging to the shared partition \mathcal{X}_{sh}^U of the search space \mathcal{X}^U . As per the evolutionary operators, we utilize, at no loss in generality, those prescribed in [22]: Simulated Binary Crossover (SBX) and polynomial mutation.

Once all details of the proposed framework have been given, we are in a position to assess its performance via computer simulation. The next section elaborates and discusses on the results obtained therefrom.

IV. EXPERIMENTS AND RESULTS

We proceed by evaluating the performance of the proposed MFEA-based framework when evolving multiple DQL models simultaneously. To this end, we structure the experiments and discussion to provide an informed answer to several research questions (RQ):

- RQ1: Can MFEA optimize DQL models individually? Do they efficiently solve the considered RL tasks?
- RQ2: Can MFEA optimize DQL models simultaneously, exploiting their synergies to achieve a better performance?
- RQ3: Can we quantify the intra- and inter-environment knowledge transfer between RL tasks realized by MFEA?

To elaborate on these questions, we first establish the $N = 3$ RL environments $\{E_n\}_{n=1}^3$ involved in the experimentation. From the set of environments provided by OpenAI Gym Library we select two environments with a discrete action space ($E_1 = \text{cartpole}$ and $E_2 = \text{acrobot}$), and a third environment with continuous action space ($E_3 = \text{pendulum}$). The configuration β of these environments has been manipulated in order to yield a total of $K = 12$ tasks $\{T_k\}_{k=1}^{12}$, which we will hereafter refer to as $\text{environment}(\alpha)$, with $\alpha \in \{A, B, C, D\}$ indicating the configuration set. Environments and configurations considered in the experiments are summarized in Table I.

TABLE I
ENVIRONMENTS AND CONFIGURATIONS OF THE EXPERIMENTS

Environments	Configuration	$\alpha = A$	$\alpha = B$	$\alpha = C$	$\alpha = D$
cartpole(α)	Pole length	0.5	0.6	0.7	0.4
acrobot(α)	Joints' length	1	1.2	1.4	1.6
pendulum(α)	Max speed/Max torque	8 / 2.0	6 / 2.0	10 / 2.0	8 / 2.5

In order to solve the aforementioned tasks, the Deep Neural Networks $\{M_{\theta_k}^k\}_{k=1}^K$ used for DQL are set with the same number of layers, parameters and overall architecture. Figure 3 (next page) depicts this common neural architecture, as well as the number of parameters on each layer. We further set $L_{sh} = 3$ shared layers among the tasks, yielding a unified search space \mathcal{X}^U with dimensionality $|\mathcal{X}^U| = 962$ decision variables (weights and biases) as per Expression (3). As for the MFEA, it is relevant to emphasize that its configuration is kept fixed for all experiments: a population size of $P = 100$

candidates, random initialization, assortative mating operator controlled by threshold $\varphi = 0.3$ (see Algorithm 2), and stopping criteria given by a total of 60 generations (namely, $6 \cdot 10^3$ fitness evaluations). Furthermore, the fitness of candidates is computed as the average of the final reward achieved by the candidate over $N_{tst} = 50$ episodes of the corresponding environment.

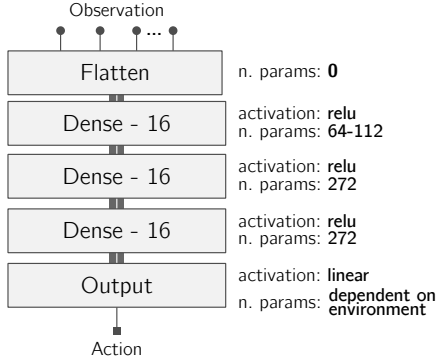


Fig. 3. Deep Neural Network architecture and number of parameters of every layer in charge of solving the environments under consideration.

To account for the statistical variability of the results, 5 independent runs of every experiment have been executed, reporting on performance statistics averaged over such runs. Each model, once optimized by the proposed framework, is tested on 250 episodes of its corresponding task, from which the performance statistics mentioned previously are computed. All the tests conducted in this experimentation have been run on an Intel(R) Xeon W-2123 processor running at 3.6 GHz with 32 Gb RAM. The source code producing the results presented in what follows has been made available at [23].

RQ1: Can MFEA optimize DQL models individually? Do they efficiently solve the RL tasks?

First, MFEA is tested on each of the tasks individually, i.e. by just considering one single task ($K = 1$). This is expected to yield the best models MFEA can achieve, setting the highest reward that can be reached on every task. Figure 4 exemplifies the evolution of the average reward attained by MFEA as a function of the number of evaluations for tasks $\text{cartpole}(B)$, $\text{acrobot}(B)$ and $\text{pendulum}(B)$. The shaded area represents the $\pm std$ range centered on the mean (bold line), computed over 5 independent runs and 250 test episodes per every run.

Some interesting aspects can be drawn from these plots: to begin with, the evolutionary operators defined in MFEA succeed at converging fast and effectively towards well-performing models, characterized by the mean and standard deviation of their reward shown at the right bottom corner of each plot. In terms of convergence, $\text{acrobot}(B)$ and $\text{cartpole}(B)$ are able to reach the highest performance in approximately 10 generations, whereas $\text{pendulum}(B)$ converges in around 20 generations. This relatively worse convergence featured by $\text{pendulum}(\alpha)$, which was indeed

observed $\forall \alpha \in \{A, B, C, D\}$, will have a strong implication in the results obtained to gain insights on RQ3.

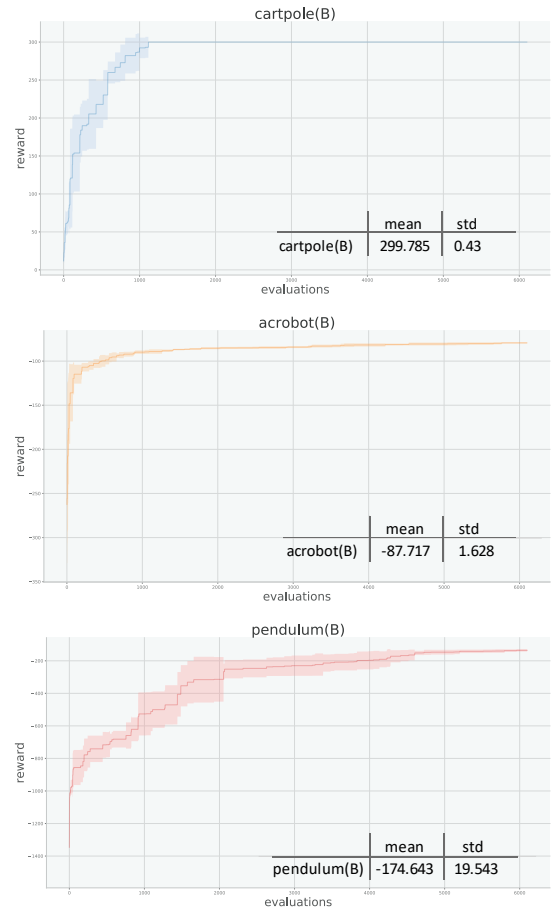


Fig. 4. Evolution of MFEA when individually evolving DQL models aimed to solve tasks $\text{cartpole}(B)$, $\text{acrobot}(B)$ and $\text{pendulum}(B)$. Mean and standard deviation values of the evolved model tested on 500 episodes of their correspondent environment is shown.

RQ2: Can MFEA optimize DQL models simultaneously, exploiting their synergies to achieve a better performance?

We proceed by inspecting the multi-environment evolving capabilities of MFEA. A single population of solutions defined on the unified search space \mathcal{X}^U is shared between all tasks, while the evolution of each individual task is performed with the addition of the intrinsic knowledge transfer endowed by MFEA. The main goal of this experiment is to check whether MFEA is a good solver for multiple environments, giving quantitative metrics of the quality (reward) achieved by models evolved via a single search process. Therefore, we retrieve the reward statistics obtained for every task when the corresponding DQL model is evolved by MFEA in isolation with respect to the other tasks (RQ1). We then compare these scores to those obtained when all tasks corresponding to a given environment are jointly evolved, potentially promoting intra-environment knowledge transfer. Since the number of fitness evaluations is kept fixed to $60 \cdot 10^3$ in all cases, the difference among such scores will evince the performance

degradation when all models are jointly optimized via MFEA.

TABLE II

AVERAGE AND STANDARD DEVIATION OF REWARDS OBTAINED BY MFEA ON DIFFERENT ENVIRONMENTS WHEN OPTIMIZING MODELS SEPARATELY OR SIMULTANEOUSLY (INTRA-ENVIRONMENT TRANSFER)

		Separately		Simultaneously	
		Mean	Std	Mean	Std
Exp. 1 cartpole	cartpole(A)	299.95	0.10	261.944	48.094
	cartpole(B)	299.79	0.43	280.098	38.178
	cartpole(C)	299.97	0.05	274.072	36.804
	cartpole(D)	300.00	0.00	287.082	15.272
Exp. 2 acrobot	acrobot(A)	-77.75	1.13	-82.732	2.651
	acrobot(B)	-87.72	1.63	-95.412	5.634
	acrobot(C)	-96.91	3.51	-101.836	4.588
	acrobot(D)	-103.07	2.32	-112.96	4.861
Exp. 3 pendulum	pendulum(A)	-270.39	99.44	-800.915	171.822
	pendulum(B)	-174.63	19.54	-562.826	138.927
	pendulum(C)	-176.83	19.65	-722.138	220.063
	pendulum(D)	-161.58	9.81	-516.174	37.71

Table II lists the results obtained for addressing RQ2, comprising three experiments (one per environment). By analyzing them, it can be claimed that MFEA allows evolving multiple DQL models efficiently by using the designed unified search space. When optimizing for the `cartpole` environment (Exp. 1), it can be seen that the DQL models evolved can perform competitively with their individually evolved counterparts, even if being granted the same computational budget. Likewise, the `acrobot` environment (Exp. 2) does not undergo a strong degradation either, reaching rewards close to those of the individually evolved models. However, in `pendulum` (Exp. 3) notably worse results were obtained. This fact unveils that the complexity of the environment to be solved directly impacts on the quality of the multifactorial evolution. The spotlight must be placed on the reasons for this degradation, reason for which an study on the transferability of knowledge among tasks (RQ3) is addressed in the next section.

RQ3: Can we quantify the intra- and inter-environment knowledge transfer among RL tasks promoted by MFEA?

We end our discussion on the performance of the proposed framework by delving into the genetic knowledge transfer of MFEA. In general, Transfer Learning is applied among Deep Neural Networks by copying weights and biases from a pre-trained network in a origin task T_k to a network aimed to solve a target task $T_{k'}$. In MFEA, however, the knowledge transfer is done throughout the evolution process, sharing the knowledge acquired by the population and embedded in the genotype of the search space \mathcal{X}^U . To undertake this knowledge transferability study, we set MFEA to evolve simultaneously $K = 9$ tasks corresponding to 3 different configurations (namely, A , B and C as per Table I) of each of the $N = 3$ environments considered in our experiments. Once this experiment is designed and run, we record similar final reward statistics to those reported for RQ1 and RQ2.

Results are listed in Table III, which must be analyzed by comparing them to those in Table II. Relatively lower

TABLE III
AVERAGE AND STANDARD DEVIATION OF REWARDS OBTAINED BY MFEA WHEN EVOLVING JOINTLY DQL MODELS AIMED AT DIVERSE RL TASKS (INTRA- AND INTER-ENVIRONMENT TRANSFER)

	$\alpha = A$		$\alpha = B$		$\alpha = C$	
	Mean	Std	Mean	Std	Mean	Std
cartpole(α)	229.21	57.6	273.87	32.25	218.82	57.21
acrobot(α)	-84.52	3.11	-95.86	2.47	-116.221	5.79
pendulum(α)	-1086.02	69.52	-580.34	72.80	-721.88	144.43

average reward values are reached for all the tasks when evolved jointly, with gaps being narrower for `acrobot` tasks. Despite this worse performance, it is important to bear in mind that in this case, more – and more diverse – tasks are evolved simultaneously with the same evaluation budget and population size. This demonstrates the potential of MFEA to evolve DQL models for complicated RL tasks at the same time. On the other hand, it provides evidence that the MFEA operators devised to promote genetic transfer among tasks do not outweigh the limited computational budget allocated for all RL problems. It is appropriate to examine the reasons for this lack of effectiveness, and discriminate whether the joint evolution of tasks defined in these three environments undergoes negative information transfer among such tasks that cannot be evaded by MFEA.

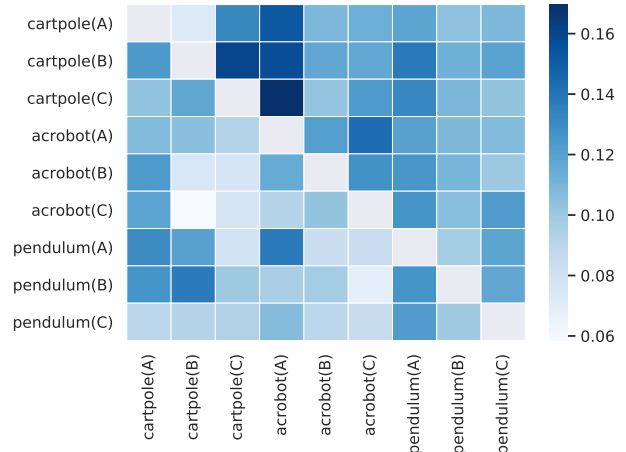


Fig. 5. Effective crossover percentage between tasks, denoting the percentage of times that a crossover with the task denoted in the row has been effective for the task denoted in the column.

To this end, we check the times an *effective crossover* has occurred among tasks, referring as such to the crossover event in which a candidate with skill factor equal to task k , when recombined with another candidate with skill factor k' , improves the performance of the latter in task $T_{k'}$. The heat map depicted in Figure 5 indicates the ratio of effective crossovers for every pair of tasks as per row and column labels. For instance, a highly intense color in the cell with coordinates (row,column)= $(T_k, T_{k'})$ indicates that the crossover operation between an individual with skill factor k , when mating with an individual with skill factor equal to k' , has positively contributed to solving $T_{k'}$. Interesting relationships emerge from this plot, such as a remarkably positive knowledge trans-

fer between `cartpole(C)` and `acrobot(A)`, followed by lower yet equally notable contributions from `cartpole(B)` and `cartpole(A)`. On the contrary, tasks whose average reward as per Table III was found to be poor seem not to contribute significantly to the convergence of tasks defined on other environments (i.e. low inter-environment transfer), as can be read for e.g. `pendulum(C)`. However, for this latter example a higher intra-environment transfer is noted, with `pendulum(B)` having received a positive push recurrently when mating with `pendulum(C)`.

V. CONCLUSIONS AND FUTURE RESEARCH

This work has presented the application of Multifactorial Optimization to simultaneously optimize multiple DQL models devised for a set of interrelated RL tasks. The goal is to evolve jointly the weights and biases of the Deep Neural Networks inside the aforementioned DQL models by formulating it as a single optimization problem defined on an unified search space, whose design is inspired by the usual way of performing Transfer Learning among Deep Neural Networks. A framework relying on the so-called Multifactorial Evolutionary Algorithm (MFEA) has been designed to efficiently explore this unified search space, incorporating further operators to probabilistically exploit possible synergies among the tasks under consideration during the search process.

We have empirically gauged the performance of MFEA to solve multiple RL scenarios at the same time. MFEA has been shown to perform well for this scenario, evolving up to nine tasks by using a unique evolution process and just one population, which is shared between all the environments/tasks. We have also evaluated the effectiveness of the knowledge transfer when evolving DQL models, with interesting insights between the complexity of the RL task and its contribution in terms of knowledge transfer to the rest of the tasks.

Future work will be devoted to the design of innovative knowledge sharing strategies, as well as new encoding schemes that extend the capabilities of the devised framework, i.e. evolution of larger networks and more complex environments. Efforts will be also invested towards optimizing DQL tasks by MFEA in an on-line fashion, i.e. while partially optimized DQL models interact with the environment.

ACKNOWLEDGMENTS

Aritz D. Martinez, Eneko Osaba and Javier Del Ser would like to thank the Basque Government for its funding support through the EMAITEK and ELKARTEK programs. Javier Del Ser receives funding support from the Consolidated Research Group MATHMODE (IT1294-19) granted by the Department of Education of the Basque Government.

REFERENCES

[1] B. Baker, I. Kanitscheider, T. Markov, Y. Wu, G. Powell, B. McGrew, and I. Mordatch, "Emergent tool use from multi-agent autotutorials," *arXiv:1909.07528*, 2019.

[2] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?," in *Advances in Neural Information Processing Systems*, pp. 3320–3328, 2014.

[3] A. Gupta, Y.-S. Ong, and L. Feng, "Insights on transfer optimization: Because experience is the best teacher," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 1, pp. 51–64, 2017.

[4] L. Feng, Y.-S. Ong, A.-H. Tan, and I. W. Tsang, "Memes as building blocks: a case study on evolutionary optimization+ transfer learning for routing problems," *Memetic Computing*, vol. 7, no. 3, pp. 159–180, 2015.

[5] Y.-W. Wen and C.-K. Ting, "Parting ways and reallocating resources in evolutionary multitasking," in *IEEE Congress on Evolutionary Computation*, pp. 2404–2411, 2017.

[6] A. Gupta and Y.-S. Ong, "Genetic transfer or population diversification? deciphering the secret ingredients of evolutionary multitask optimization," in *IEEE Symposium Series on Computational Intelligence*, pp. 1–7, 2016.

[7] A. Gupta, Y.-S. Ong, and L. Feng, "Multifactorial evolution: toward evolutionary multitasking," *IEEE Transactions on Evolutionary Computation*, vol. 20, no. 3, pp. 343–357, 2015.

[8] Y.-S. Ong and A. Gupta, "Evolutionary multitasking: A computer science view of cognitive multitasking," *Cognitive Computation*, vol. 8, no. 2, pp. 125–142, 2016.

[9] J. C. Bean, "Genetic algorithms and random keys for sequencing and optimization," *ORSA Journal on Computing*, vol. 6, no. 2, pp. 154–160, 1994.

[10] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv:1312.5602*, 2013.

[11] F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune, "Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning," *arXiv:1712.06567*, 2017.

[12] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Processing Magazine*, vol. 34, pp. 26–38, Nov 2017.

[13] C. Igel, "Neuroevolution for reinforcement learning using evolution strategies," in *IEEE Congress on Evolutionary Computation*, vol. 4, pp. 2588–2595, 2003.

[14] S. Khadka and K. Tumer, "Evolutionary reinforcement learning," *arXiv:1805.07917*, 2018.

[15] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, "Evolution strategies as a scalable alternative to reinforcement learning," *arXiv:1703.03864*, 2017.

[16] E. Conti, V. Madhavan, F. P. Such, J. Lehman, K. Stanley, and J. Clune, "Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents," in *Advances in Neural Information Processing Systems*, pp. 5027–5038, 2018.

[17] J. Koutník, J. Schmidhuber, and F. Gomez, "Evolving deep unsupervised convolutional networks for vision-based reinforcement learning," in *Annual Conference on Genetic and Evolutionary Computation*, pp. 541–548, 2014.

[18] R. Glatt, F. L. Da Silva, and A. H. R. Costa, "Towards knowledge transfer in deep reinforcement learning," in *Brazilian Conference on Intelligent Systems*, pp. 91–96, 2016.

[19] M. E. Taylor and P. Stone, "Transfer learning for reinforcement learning domains: A survey," *Journal of Machine Learning Research*, vol. 10, no. Jul, pp. 1633–1685, 2009.

[20] L. Shao, F. Zhu, and X. Li, "Transfer learning for visual categorization: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 5, pp. 1019–1034, 2014.

[21] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "OpenAI gym," *arXiv:1606.01540*, 2016.

[22] B. Da, Y.-S. Ong, L. Feng, A. K. Qin, A. Gupta, Z. Zhu, C.-K. Ting, K. Tang, and X. Yao, "Evolutionary multitasking for single-objective continuous optimization: Benchmark problems, performance metric, and baseline results," *arXiv:1706.03470*, 2017.

[23] A. D. Martinez, E. Osaba, and J. Del Ser, "Implementation of MFEA for Deep Q Learning." <https://git.code.tecnalia.com/aritz.martinez/dql-mfea>, 2019.