

EvoPruneDeepTL: An evolutionary pruning model for transfer learning based deep neural networks

Javier Poyatos^a, Daniel Molina^{a,*}, Aritz D. Martinez^b, Javier Del Ser^{b,c}, Francisco Herrera^{a,d}

^a Department of Computer Science and Artificial Intelligence, Andalusian Research Institute in Data Science and Computational Intelligence (DaSCI), University of Granada, Granada, 18071, Spain

^b TECNALIA, Basque Research & Technology Alliance (BRTA), Derio, 48160, Spain

^c University of the Basque Country (UPV/EHU), Bilbao, 48013, Spain

^d Faculty of Computing and Information Technology, King Abdulaziz University, Jeddah, 21589, Saudi Arabia

ARTICLE INFO

Article history:

Received 6 February 2022
Received in revised form 27 September 2022
Accepted 11 October 2022
Available online 4 November 2022

Keywords:

Deep learning
Evolutionary algorithms
Pruning
Feature selection
Transfer learning

ABSTRACT

In recent years, Deep Learning models have shown a great performance in complex optimization problems. They generally require large training datasets, which is a limitation in most practical cases. Transfer learning allows importing the first layers of a pre-trained architecture and connecting them to fully-connected layers to adapt them to a new problem. Consequently, the configuration of these layers becomes crucial for the performance of the model. Unfortunately, the optimization of these models is usually a computationally demanding task. One strategy to optimize Deep Learning models is the pruning scheme. Pruning methods are focused on reducing the complexity of the network, assuming an expected performance penalty of the model once pruned. However, the pruning could potentially be used to improve the performance, using an optimization algorithm to identify and eventually remove unnecessary connections among neurons. This work proposes EvoPruneDeepTL, an evolutionary pruning model for Transfer Learning based Deep Neural Networks which replaces the last fully-connected layers with sparse layers optimized by a genetic algorithm. Depending on its solution encoding strategy, our proposed model can either perform optimized pruning or feature selection over the densely connected part of the neural network. We carry out different experiments with several datasets to assess the benefits of our proposal. Results show the contribution of EvoPruneDeepTL and feature selection to the overall computational efficiency of the network as a result of the optimization process. In particular, the accuracy is improved, reducing at the same time the number of active neurons in the final layers.

© 2022 Published by Elsevier Ltd.

1. Introduction

Deep Learning (DL) (Goodfellow et al., 2016) is one of the most attractive research areas in machine learning in recent times, due to the great results offered by such models in a plethora of applications. DL architectures are successfully used in many problems, like audio classification (Lee et al., 2009, December), audio recognition (Noda et al., 2015), object detection (Zhou et al., 2017, May), image classification for medical analysis (Muhammad et al., 2021) or vehicular perception (Muhammad et al., 2020).

Convolutional Neural Networks (CNNs) (Lecun et al., 1998) constitute the state-of-the art in image classification (Sultana

et al., 2018, November). CNNs include two parts, the first part is actually a feature extractor based on convolution and pooling operations. The second part usually contains one or more fully connected layers. In these fully-connected layers, the neuron of each layer is connected to all the neurons of the previous layer, which generates a large number of weights to be trained. The design of an appropriate network for each problem is a requirement in order to obtain a good performance. The training process of a DL architecture is frequently time-consuming. Complexity reduction maintaining the performance is an important challenge in DL, currently attracting significant attention in the community. Transfer Learning (TL) (Weiss et al., 2016) is usually considered the alternative. It is very common to use a DL model with fixed and pre-trained weights in the convolutional layers with a dataset (like ImageNet Krizhevsky et al., 2012, December) and then add and train several layers, named fully-connected layers, to adapt the network to a different classification problem

* Corresponding author.

E-mail addresses: jpoyatosamador@ugr.es (J. Poyatos), dmolina@decsai.ugr.es (D. Molina), aritz.martinez@tecnalia.com (A.D. Martinez), javier.delser@tecnalia.com (J. Del Ser), herrera@decsai.ugr.es (F. Herrera).

([dataset] Gómez-Ríos et al., 2019; Khan et al., 2019; Shin et al., 2016).

The architecture of fully-connected layers used for the problem is a critical decision, and its design is still an open issue in terms of the number of layers and neurons per layer (Liu, Wang, Liu et al., 2017). There are general guidelines based on the experience working with these layers, rather than rules to follow for the configuration of them. Therefore, any kind of optimization in them could provide a benefit in terms of model complexity and performance. The pruning approaches follow the key idea of reducing the complexity of the model, which creates new networks with less computational cost for training. This idea is followed in Frankle and Carbin (2019, May), which also shows that, in the end, the accuracy can also improve as a result of pruning.

Pruning is interpreted as removing unnecessary connections from the model, but learning which connections are the fittest to improve the performance of the model is the key point. In fact, the selection of the best *features* for the problem is known as Feature Selection (FS) (Iguyon & Elisseeff, 2003). In our case, TL allows the extraction of the features of the input data of the DL model. These features are the input of the fully-connected layers that will be trained and, as a result of that, will largely affect the performance of the network. Nonetheless, in many cases, the problem that is formulated to learn these features is usually different, sometimes more complex, than the one at hand and, therefore, not all the learned patterns would be required. For that reason, FS gives rise to an interesting option to select and retain the subset of all features that lead to an improved performance of the model (Yildirim et al., 2018).

In pruning scenarios, the main aim of most of the traditional pruning techniques mainly aim at reducing the number of trainable parameters of the network, at the cost of a lower performance. They seek to control the performance degradation resulting from the process, but it is not their priority. Furthermore, they locally optimize parts of the network rather than searching for globally optimal pruning policies, yielding usually suboptimal pruned subnetworks with a lower performance. Another disadvantage of these pruning proposals is the fact that, as the pruning affect all layers, the complete network must be trained again, hence obtaining no advantages from the TL process. It could be useful to have a pruning technique that prioritizes results over complexity reduction, targeting a global performance improvement of the network while reducing its complexity.

Transforming the fully-connected layers into a sparse representation, in which each connection could be active or inactive, could be used to prune neural networks. Following this approach, both pruning and FS can be seen as optimization problems, in which the target is to obtain the active set of connections that produce the best performance. This optimization problem can be globally tackled by optimization algorithms like Evolutionary Algorithms (Back et al., 1997) (EAs). They have been successfully applied to many complex optimization problems. Even though they cannot guarantee the achievement of the optimum for the problem at hand, they obtain good results with limited resources and reasonable processing time. Another advantage is their versatility: several of them, like genetic algorithms (Goldberg, 1989) (GAs) allow optimizing solutions with different representations (Chambers, 2000). The spectrum of problems in which EAs can be used is very wide. EAs have been traditionally applied to optimize neural networks (Iba, 2018), but their usage in DL networks to improve DL networks (Martinez et al., 2021), to train them (Mohapatra et al., 2022), and to create new DL networks from scratch (Elsken et al., 2019b) is more recent. The use of EA's is mainly oriented towards optimizing a complete network. However, in this paper, our aim is to adapt the fully-connected layers

(the only trained for the problem to solve using TL) to improve the accuracy in the predictions, together with the complexity reduction. Our main hypothesis is the convenience of use of EAs to prune the fully-connected layers via a sparse representation.

We propose an evolutionary pruning model based on TL for deep neural networks, Evolutionary Pruning for Deep Transfer Learning (EvoPruneDeepTL). EvoPruneDeepTL can be applied to a DL model that resorts to TL to tackle a new task. EvoPruneDeepTL combines sparse layers and EA, consequently, neurons in such layers are pruned to adapt their sparsity pattern to the addressed problem. EvoPruneDeepTL is able to efficiently explore the neuron search space (to discover coarsely grained solutions) or, alternatively, in the connection search domain (fine-grained solutions).

An important aspect to analyze in EvoPruneDeepTL is that one of its solution encoding schemes effectively leads to a feature selection mechanism, in which we deactivate the extracted features and the EA evolves these features to learn which ones fit best as predictors for the given problem.

EvoPruneDeepTL' goals include flexibility and adaptability. EvoPruneDeepTL has been designed to be flexible, and the automatic configuration of the network can be applied to different pre-trained networks, used as feature extractors, and different fully-connected layers. This make our proposal capable of tackling different problems. The optimization in connections that EvoPruneDeepTL generates allows the model to be adaptable to the specific dataset to be modeled. Thus, when the dataset suffers a change, the resultant configuration will also be adapted to the new circumstances.

To assess the performance of EvoPruneDeepTL, we have conducted an extensive experimentation that leads to several valuable insights. To begin with, experimental results showcase the behavior and effectiveness of EvoPruneDeepTL in terms of precision and in terms of reduction of the complexity of the network. Thanks to the flexibility of EvoPruneDeepTL, it is applied to perform either pruning or FS. Both cases improve the accuracy of the network when the comparison is made against reference models and CNN pruning methods from the literature. Moreover, in most cases, the FS scheme achieves a better performance than the pruning scheme in terms of the accuracy of the network. Furthermore, the network pruned by the FS scheme also achieves a significantly reduced number of connections in its fully connected part, contributing to the computational efficiency of the network. We have also included several experiments showing the flexibility of the model, both changing the feature extractor and showing how changes in the dataset implies a modification in the final configuration obtained by EvoPruneDeepTL. In short, this extensive experimentation is used to provide answer to the following six questions as the thread running through this experimental study:

- (RQ1) Which is the performance of EvoPruneDeepTL against fully-connected models?
- (RQ2) Which would be better, to remove neurons or connections?
- (RQ3) Which is the performance of EvoPruneDeepTL when compared to other efficient pruning methods?
- (RQ4) Which would be better, the use of pruning of fully-connected layers or Feature Selection?
- (RQ5) How does EvoPruneDeepTL perform when applied to different pre-trained networks?
- (RQ6) Can EvoPruneDeepTL adapt efficiently their pruned knowledge to changes in the modeling task, showing robustness?

The rest of the article is structured as follows: Section 2 exposes related work to our proposal present in the literature. Section 3 shows the details of the proposed EvoPruneDeepTL model. Section 4 presents our experimental framework. In Section 5, we show and discuss the EvoPruneDeepTL's results of the experiments of pruning, feature selection and against efficient CNN pruning methods of the literature. Moreover, EvoPruneDeepTL is tested with different extractor features and with different variations of datasets in this section. Section 6 follows by summarizing the advantages and drawbacks of our proposal when compared to other pruning approaches. Finally, Section 7 draws the main conclusions stemming from our work, and outlines future work departing from our findings.

2. Related work

The purpose of this section is to make a brief review of contributions to the literature that link to the key elements of our study: Transfer Learning (Section 2.1), Neural Architecture Search (Section 2.2), CNN pruning (Section 2.3), Evolutionary Algorithms (Section 2.4) and Feature Selection with Deep Learning (Section 2.5).

2.1. Transfer learning

TL (Pan & Yang, 2010) is a DL mechanism encompassing a broad family of techniques (Tan et al., 2018, October). Arguably, the most straightforward method when dealing with neural networks is *Network-based deep transfer learning*, in which a previous network structure with pre-trained parameters in a similar problem is used. It offers good results by the behavior of DL models, in which first layers detect useful features on the images, and later layers strongly depend on the chosen dataset and task. As finding these standard features on the first layers seems very common regardless of the natural image datasets, its trained values can be used for different problems (Yosinski et al., 2014, December). Training DL models from scratch is usually time-consuming due to the great amount of data in most cases. TL gives some benefits which make it a good option for DL: reduction of time needed for training (Sa et al., 2016), better performance of the model and less need of data.

TL has been applied to several real-world applications, such as sound detection (Jung et al., 2019, May) or coral reef classification (Gómez-Ríos et al., 2019). Moreover, in Tajbakhsh et al. (2016) two different approaches for TL are discussed: fine-tuning or full training. They demonstrated that, for medical reasons, a pre-trained CNN with adequate fine-tuning performed better in terms of accuracy than a CNN trained from scratch. Another approach of TL is presented in Mehdiipour Ghazi et al. (2017), in which an optimization of TL parameters for plant identification is proposed.

There are different deep neural networks proposed in the literature. One of the most popular is ResNet, which uses residual learning to improve the training process, obtaining better performance than other models (He et al., 2016, June). ResNet models are characterized by the use of deeper neural networks without loss of information due to their architecture. Different ResNet models with TL have been used in several applications (Scott et al., 2017), such as medical classification like pulmonary nodule (Nibali et al., 2017) and diabetic retinopathy classification (Wan et al., 2018). Moreover, other networks have shown great performance when used with TL, such as DenseNet (Huang et al., 2017, July) and VGG (Simonyan & Zisserman, 2015, May). An example of DenseNet with TL is presented in Aneja and Aneja (2019, July), which shows that this network architecture is able to achieve a great result for the task at

hand when combined with TL. Lastly, VGG has also shown an outstanding performance when it is used in combination with TL. An example is presented in Wen et al. (2019, May) in which a pre-trained VGG-19 network is used to solve a fault diagnosis problem.

2.2. Neural architecture search

The appropriate design of a neural network is a key point to solve DL problems. Nevertheless, finding the best architecture that optimally fits the data and, as a result of that, gives the best outcome for the problem is extremely difficult. Recently, the term Neural Architecture Search, NAS, has obtained a great importance in this field. The objective of NAS is the automatic search for the best design of a NN to solve the problem at hand.

The first work in this field is presented in Stanley and Mikulainen (2002) in the beginning of this century, in which they demonstrate the effectiveness of a GA to evolve topologies of NN.

In Zoph et al. (2018, June), the authors design the NASNet architecture, a new search space to look for the best topology for the tackled problem. Moreover, in Liu, Zoph et al. (2018, September) the authors propose to search for structures in increasing order of their level of complexity, while learning a surrogate model to guide the search through structure space.

NAS methods usually rely on Reinforcement Learning, RL, and EA, like (Zoph et al., 2018, June) or (Liu, Simonyan et al., 2018, April), in which the authors explore the search space using a hierarchical genetic representation. Another example of RL for NAS is shown in Kokiopoulou et al. (2020, August). The authors propose a novel method that, by sharing information on multiple tasks, is able to efficiently search for architectures.

NAS can also be viewed as a multi-objective problem. Among these methods, one of them is presented in Elsken et al. (2019a, May), in which the authors propose a multi-objective for NAS that allows approximating the entire Pareto-front of architectures. Another example is Neural Architecture Transfer (Lu et al., 2021), that allows to overcome a common limitation of NAS, that is requiring one complete search for each deployment specification of hardware or objective. They use an integrated online transfer learning and a many-objective evolutionary search procedure.

Recently, one of the most well-known multi-objective EA, NSGA-II, has been used for NAS (Lu et al., 2019, July), called NSGA-Net. This novel proposal looks for the best architecture through a three-step search: an initialization step from hand-crafted architectures, an exploration step that performs the EA operators to create new architectures, and finally an exploitation step that utilizes the knowledge stored in the history of all the evaluated architectures in the form of a Bayesian Network.

Lastly, there are more advanced techniques of NAS and EA given by Real et al. (2019, January), in which a new model for evolving a classifier is presented, and by Real et al. (2020, July), in which the authors propose AutoML-Zero, an evolutionary search to build a model from scratch (with low-level primitives for feature combination and neuron training) which is able to get a great performance over the addressed problem.

2.3. CNN pruning

The main reason to optimize the architecture of a deep neural network is to reduce its complexity. That reduction can be done in different ways (Long, Ben and Liu, 2019). One of them is by designing compact models from scratch instead of resorting to architectures comprising multiple layers. Another strategy is via weights-sharing (Ullrich et al., 2017, April). An alternative method to reduce the complexity of DL models is low-rank factorization (Long, Ben, Zeng et al., 2019), based on a matrix

decomposition to convolutional layers to estimate parameters. However, one of the most popular is Network Pruning. The objective of pruning is to remove unnecessary parameters from a neural network, so that they do not participate during training and/or inference. It can be done in the convolutional phase on the channels, kernels and weights or even in the fully connected phase on the neurons. In [Masson et al. \(2021\)](#) they show a classification of pruning methods for channels. They categorize the pruning methods for channel reduction, and they also specify the criteria used to select these channels: based on weights or based on feature maps.

We have seen that network pruning has achieved a great importance in the literature as many researchers have applied different techniques to simplify a CNN using a pruning scheme. In [Liu et al. \(2019, May\)](#) they classify the pruning methods in unstructured and structured pruning, and make a review of all the state-of-art structured pruning methods. Unstructured pruning methods remove weights without following any order. For the structured methods, there are some rules or even constraints which define how the pruning is done ([Anwar et al., 2017](#)). Typically, the pruned layers appertain to the convolutional phase ([Luo et al., 2017, October](#)). In our proposal, we instead apply a structured pruning scheme to the fully-connected layers.

Among pruning methods, the value-based *weight pruning* ([Han et al., 2015, December](#)) and *neuron pruning* ([Srinivas & Babu, 2015, September](#)) have arisen as the most used, particularly due to their simplicity. The logic behind this pruning methods is straightforward: a certain amount (%) of the weights or neurons that contribute less to the final trained model are removed from the architecture. This makes the network quicker to perform inference and endows it with better generalization capabilities. However, multiple pruning and retraining steps demonstrated that it is possible to recover fully or partially the knowledge lost in the pruning phase. Further along the series of pruning approaches published to date, *Polynomial Decay* ([Zhu & Gupta, 2018, April](#)) is a scheduled pruning method that considers that a higher amount of weights can be pruned in early stages of pruning, while systematically less amount of weights should be pruned in late stages. Between pruning steps, the network is retrained for some epochs. An implementation of the discussed methods can be found for Tensorflow.¹

Pruning a CNN model reduces its complexity, but sometimes leads to a decrease of the performance of the model, although there are some proposals that reduce the complexity of the model with no loss of accuracy ([Han et al., 2016, May](#)).

Pruning a neural network can be conceived as an optimization method in which we start from the original vector, and connections/neurons are decision variables whose value is evolved towards optimizing a given objective. In this context of evolution of neural networks, evolutionary algorithms for evolving DL architectures have been applied ([Iba, 2018](#)). While this combination of EA and DL models seemed to be a great scheme for the optimization of DL models, especially for CNN network, the optimization of DL models is still an open problem ([Liu, Wang, Liu et al., 2017](#)). Many proposals have been published about this problem like in [Martinez et al. \(2021\)](#), where they make a review of proposals using EAs for optimizing DL models, prescribing challenges and future trends to effectively leverage the synergy between these two areas.

Researchers have presented a great variety of proposals about the optimization of DL models using EA, most commonly for CNN. In [Martín et al. \(2018\)](#), the authors developed EvoDeep, an EA with specific mutation and crossover operators to automatically

create DL models from scratch. Moreover, in [Real et al. \(2017, August\)](#) a novel evolution approach to evolve CNN models using a GA was proposed. Another example of the optimization of CNN was developed in [Assunção et al. \(2019\)](#), in which a GA was presented for the optimization of the topology and parameters of the CNN.

In our proposal, we improve the performance of the models using a TL approach to extract the features of the images and apply a reduction of the fully-connected layer using a GA to optimize a sparse layer.

2.4. Evolutionary algorithms for CNN pruning

In the previous section, several works of CNN pruning have been presented, but none of them use an EA to prune. In this section, we mention some studies present in the literature which have used an EA in order to prune a CNN model. To begin with, in [Liu, Wang and Qiao \(2017, February\)](#), the authors propose a sparse approach to reduce CNN complexity. EAs are also a good way to prune CNN. In [Mantzaris et al. \(2011\)](#), a first attempt of pruning and GA is proposed for a medical application. They use a GA to search for redundancy factors in a neural network. Moreover, in [Samala et al. \(2018\)](#) another EA is presented to prune deep CNN for breast cancer diagnosis in digital breast tomosynthesis. A combined approach of EA and sparse is proposed by [Wang et al. \(2020\)](#), in which a GA and sparse learning are applied to a scheme of network channel pruning in the convolutional scheme of the CNN. For pruning CNN, not only GAs but also other algorithms are used, like Differential Evolution (DE). In [Salehinejad and Valaee \(2021\)](#) the authors propose to use a Differential Evolution algorithm to prune the convolutional phase and the fully-connected phase of some Deep CNN, obtaining a reduction of the model but a small decrease of its performance.

However, all previous works are focused on reducing the complexity, using the EA to reduce the accuracy loss of the pruned network. Also, many of them try to reduce the whole model, changing the complete architecture and making the pre-trained values unusable. The re-training of the network may be a time-consuming task, so we assume that TL is useful in this context. We therefore maintain the original architecture with pre-trained values. Our model focuses on improving the performance of the model by pruning connections of the fully-connected layers using a GA to evolve the connections. In this environment, the search space of the GA is narrower and a faster convergence of the algorithm may be reached.

In addition to that, in the field of neural architecture search, more advanced techniques have been developed. Among them, in Section 2.2, either ([Real et al., 2019, January](#)) and ([Real et al., 2020, July](#)) have been commented. Nonetheless, they also have a great relevance in this section. The first one evolves a classifier, whereas in the second one, the authors develop an evolutionary search to build a model from scratch.

2.5. Feature selection and deep learning

One of the advantages of using TL is reducing the required time to train a DL architecture. Nonetheless, the result of this process may lead to recognize patterns that are not useful to address the problem at hand. For that reason, once TL is applied, a FS process to obtain the best features might lead to a better performance of the neural network ([Roy et al., 2015, July](#)).

An example of this process is presented in an arrhythmia detection task addressed in [Yildirim et al. \(2018\)](#), in which the authors propose a mechanism based on feature extraction and selection to improve and ultimately obtain one of the best results

¹ https://www.tensorflow.org/model_optimization/guide/pruning, Tensorflow Pruning. Last access: 28/01/2022.

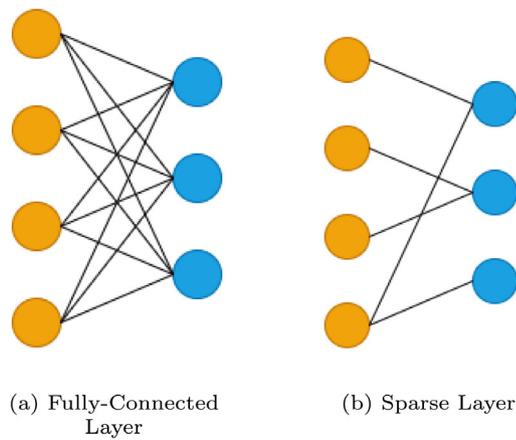


Fig. 1. Representation of both architectures.

for this problem. In relation to medical problems, the combination of FS and DL is also used in cancer diagnosis and digital breast tomosynthesis. In Samala et al. (2018) they use a TL approach and then a FS process followed by an evolution through a GA that leads to a reduced network with the same performance. Another example is described for remote sensing scene classification, in which the FS makes an impact to improve the performance of the neural network models (Zou et al., 2015), as the authors formulate the FS problem as a feature reconstruction problem. Their iterative method selects the best features to solve this problem as the discriminative features.

In our proposal, if we assume that TL is applied and we only have one fully-connected layer, then the pruning is made in relation to the extracted features of the network and, therefore, we are making a selection of the features that adjust at best to the tackled problem.

3. Evolutionary Pruning for Deep Transfer Learning

This section describes EvoPruneDeepTL, which is a model that replaces fully-connected layers with sparse layers optimized using a genetic algorithm in a TL approach. Section 3.1 gives a notion of the concept of sparse layer and the description of EvoPruneDeepTL. In Section 3.2, we define the evolutionary components of EvoPruneDeepTL. The description of the process of creating the network and how the pruning is made is shown in Section 3.3.

3.1. Global scheme of Evolutionary Pruning for Deep Transfer Learning

In a fully-connected layer, all neurons of each side are connected. Sometimes, all these connections may not be necessary, and the learning process can be reduced. For that reason, the fully-connected layer can be replaced by a sparse layer, in which some connections are eliminated.

In this work, our goal is to improve the performance of the neural network and, at the same time, to decrease the maximum number of connections or neurons. To this end, we use a sparse layer, which is composed of a subset of all connections of a fully-connected layer.

Fig. 1(a) shows the fully-connected network architecture, while Fig. 1(b) represents the sparsely connected architecture with a connection matrix of 4×3 because we have 3 classes (blue circles) and 4 neurons of the previous layer.

In this section, we discuss the basic notions of pruning and sparse layers. Moreover, the encoding strategies of

EvoPruneDeepTL are described, together with the decoding process of the chromosome encoding the pruning pattern (genotype) that yields the pruned sparse layer(s) (phenotype). Nonetheless, for the sake of a clear vision of EvoPruneDeepTL, Fig. 2 shows a diagram that exposes its general components. We next complement the detailed description provided in the following subsections with a short, albeit illustrative introduction to the key parts and overall workflow of EvoPruneDeepTL:

First, the TL process departs from a source dataset modeled by a CNN model, which adjusts the value of its coefficients by means of its learning algorithm. Then, those parameters are transferred to another network aimed to model a target task. This implicitly assumes that both tasks are correlated with each other, such that the knowledge delivered from the source to the target task via the transferred network weights can positively contribute to the learning process of the target task. These weights are kept fixed, *frozen*, in this study. Then, EvoPruneDeepTL specializes the fully-connected part of the neural network of the target task by resorting to a GA. This metaheuristic wrapper prunes unnecessary neurons of these layers driven by the improvement of a performance measure (e.g. accuracy). The outcome of the process is a pruned network with a potentially improved accuracy by virtue of an optimized pruning mask.

In this study, we propose a novel method to prune the neurons, that considers the removal of both single connections and groups of connections of the input connections of a specific neuron, as can be observed in Fig. 3. Fig. 3(b) shows a sparse layer that leads to the encoding strategy used in this work. This encoding, which is represented by the chromosome of the GA, is required to know exactly which connections are removed.

EvoPruneDeepTL model utilizes a GA designed to optimize the connections of a sparse layer. The GA takes each individual as a mask for the neural network and creates a sparse layer activating from the mask. This optimized mask gives rise to a pruned neural network suitable for the problem under consideration.

The optimization is performed using both methods, either by groups of connections or by single connections. The genome representation of each chromosome of the GA is binary-coded and represents the active *neurons* or the active *connections*. The GA evolves the configuration of the network towards its best pruned variant in terms of accuracy. Next, we describe both encoding strategies:

- **Neurons:** each gene of the chromosome represents the number of active neurons. A value 1 in position i means that the neuron i is active, and a 0 that is inactive. A non-active neuron implies that all the input connections are removed both in training and inference times. The length of the chromosome in this case is the number of neurons of the sparse model.
- **Connections:** each gene represents the connection between the layers. The interpretation of the binary values is as follows: if a gene is 1, the connection between the corresponding layers exists, otherwise, that connection does not exist. Therefore, the length of the chromosome is the maximum number of connections, noted as $D = D_1 \times D_2$, where D_1 is the number of neurons in the previous layer, and D_2 is the number of neurons in that layer.

An example of both encoding strategies is shown in Fig. 3. In both cases, the pruned connections are from the input on the layer, i.e. the right layer. The left image shows a representation of neuron-wise encoding, in which a group of neurons is selected to be active and the rest are pruned. The right image depicts how single connections are pruned.

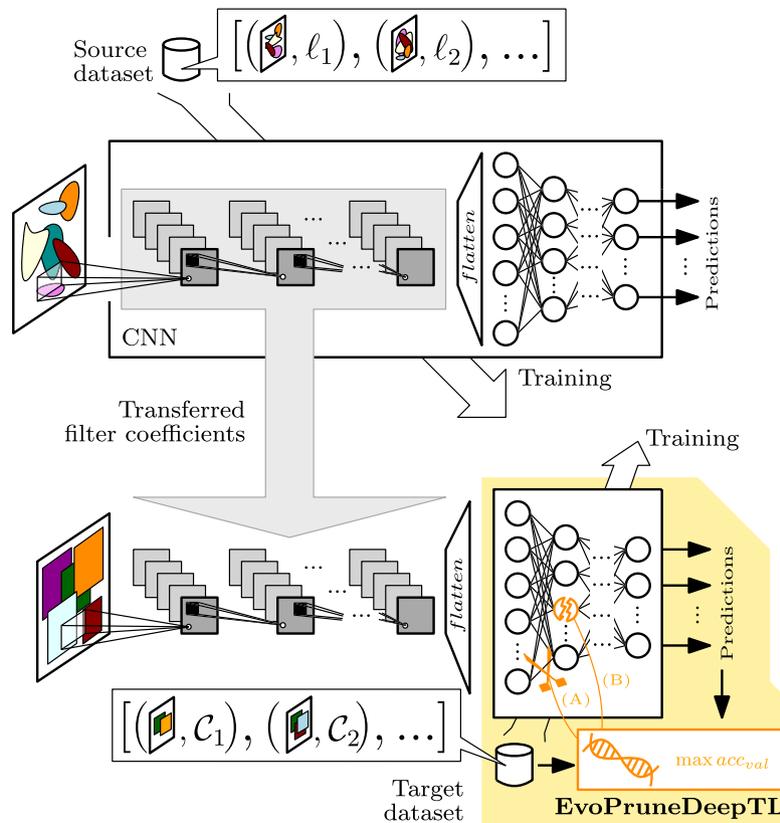


Fig. 2. Diagram of EvoPruneDeepTL.

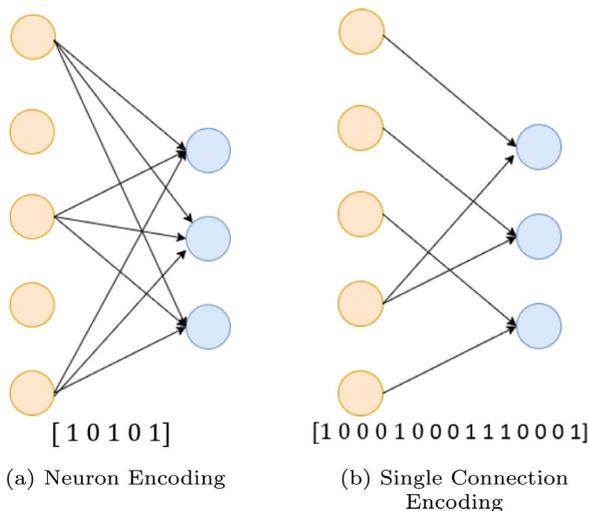


Fig. 3. Representation of encoding strategies.

3.2. Evolutionary components of EvoPruneDeepTL

In this subsection, we introduce the evolutionary components of EvoPruneDeepTL. It is a steady-state genetic algorithm, which means that two new individuals, called offsprings, are created in each generation, for the previous mentioned encoding strategies (neuron encoding vs single connection encoding, Fig. 3): in each iteration two individuals are selected and crossed, producing two offsprings that could also be mutated. The offspring candidates are introduced in the population only if they improve the worst candidates in the population, replacing them.

As previously mentioned, in EvoPruneDeepTL each chromosome is a binary array and each gene represents a connection between two layers. Each generation follows the classical scheme of selection, crossover, mutation and replacement. The best solutions found during the evolutionary search are kept in a population of individuals. Next, we describe the different components:

Selection: the implemented selection operator is Negative Assorting Mating (NAM) (Fernandes & Rosa, 2001, May). The first parent is picked uniformly at random, while the second parent is selected between three possible candidates. These candidates are also picked uniformly at random from the population. The candidate with higher Hamming distance from the first parent is chosen as the second parent, thereby ensuring that the recombined parents are diverse. This selection method allows for a higher degree of exploration of the search space.

Crossover: EvoPruneDeepTL uses the uniform crossover operator shown in Expression (1). Given two parents \mathbf{P} and \mathbf{Q} , where $\mathbf{P} = \{p_i\}_{i=1}^D$ and $\mathbf{Q} = \{q_i\}_{i=1}^D$. Then two offsprings $\mathbf{P}' = \{p'_i\}_{i=1}^D$ and $\mathbf{Q}' = \{q'_i\}_{i=1}^D$ are created following the equations:

$$p'_i = \begin{cases} p_i & \text{if } r \leq 0.5 \\ q_i & \text{otherwise} \end{cases} \quad (1)$$

$$q'_i = \begin{cases} q_i & \text{if } r \leq 0.5 \\ p_i & \text{otherwise} \end{cases}$$

where r is the realization of a continuous random variable with support over the range $[0.0, 1.0]$. This operator takes two parents \mathbf{P} and \mathbf{Q} of length D and creates two new offspring \mathbf{P}' and \mathbf{Q}' of the same size. Each new offspring is composed of the parents genes. Each gene (position of the new array) is set equal to the gene of the first or second parent. This process is repeated until the whole offspring is composed.

Mutation: EvoPruneDeepTL adopts the so-called single point mutation. A mutation probability for each individual is defined

by p_{mut} . Then, a gene of that individual is uniformly randomly selected and its bit is flipped, i.e., if the mutation is performed, then that neuron or connection changes its value, which implies that the connection or the neurons is activated or deactivated. In this operator, p_{mut} is the value that establishes the probability that a mutation is performed.

Replacement Strategy: at the end of every generation, the two offsprings resulting from the crossover and mutation operators compete against the worst two elements. As a result, the population is updated with the best two individuals among them, i.e. those whose fitness value is better. EvoPruneDeepTL maintains a pool of four individuals: two offsprings and the two worst individuals selected from the population. Then, the best two of them are in the new population. The criterion to select the best two is based on the fitness as the best of them are selected. In case of same values, the individuals with fewer active neurons/connections are those selected to be retained in the new population.

Initialization: the genes composing the individuals are initialized to 0 or 1 as per the following probabilistic condition with a p_{one} probability:

$$l_i = \begin{cases} 1 & \text{if } r \leq p_{one} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where r is the realization of a uniform continuous random variable with support $[0.0, 1.0]$.

Evaluation of individuals: the fitness value of every individual is given by the accuracy over a test dataset of the neural network pruned as per the decoded individual, and trained over the training dataset of the task at hand. Each individual, named p , is decoded to yield a sparse neural network, which we hereafter refer to $SparseNet_p$. Then this network is trained as previously commented over the train dataset, giving the $TrainedSparseNet_p$ network. Lastly, the test dataset is evaluated in this network, producing the fitness of the individual, which we call $ChildFitness_p$.

Algorithm 1 shows the pseudocode of EvoPruneDeepTL. First, we need to understand what EvoPruneDeepTL requires to start its evolution process, and what results from this process. The input of EvoPruneDeepTL is determined by:

- Dataset and task to be modeled.
- Configuration of the GA: parameters needed for the algorithm.
- Configuration of the network: parameters needed for the network.
- Feature extractor: a pre-trained neural network used for feature extraction and TL, e.g., ResNet-50 trained over Imagenet or any other available architecture alike.

The algorithm starts by initializing the individuals of the population (line 1) using the previous operator and then evaluating them (line 2). The evolutionary process is performed in lines 3–15. Two parents are selected using the NAM operator (line 4) and then the two offsprings are generated using the crossover operator (line 5). If the mutation condition is met, then mutation is performed (lines 6–8). The child population is now evaluated (lines 9–14). The evaluation is held over three steps, in which each individual is decoded (variable $SparseNet_p$ in line 10), and then the network created with its configuration is trained (named $TrainedSparseNet_p$ in line 11) using the train dataset and then evaluated (called $ChildFitness_p$ in line 12) over the test dataset. Lastly, the replacement strategy is triggered (line 15). The stopping criterion is the evaluation of a maximum number of networks.

Algorithm 1: EvoPruneDeepTL

```

Input : Dataset, configuration of the GA, configuration of
         the network and feature extractor
Output: Best pruned network
1 Initialization of individuals of the population using the
  initialization operator;
2 Evaluation of the initial population (see lines 9-14);
3 while evaluations < max_evals do
4   Parent selection using NAM operator;
5   Generate offsprings using crossover operator;
6   if rand() < p_mut then
7     Perform mutation using mutation operator;
8   end
9   for each child  $p$  in children population do
10    SparseNet $p$  ← Create sparse network using the
      decoded individual of the population;
11    TrainedSparseNet $p$  ← Train SparseNet $p$  using train
      dataset;
12    ChildFitness $p$  ← Accuracy of TrainedSparseNet $p$ 
      evaluated in test dataset;
13    evaluations+=1;
14  end
15 Replacement Strategy: child population vs worst
    individuals of population;
16 end

```

3.3. EvoPruneDeepTL network

This subsection is devised to fully understand the components associated with the networks that involve EvoPruneDeepTL. EvoPruneDeepTL stands for the usage of transfer learning, which means that the convolutional phase before the fully-connected layers is imported from other pre-trained model. Thus, the chosen CNN works as a feature extractor, i.e, obtains the main feature or characteristics for the task at hand. In our study, we have chosen ResNet-50 as feature extractor, although others can also be used, such as VGG or DenseNet. In Section 5, a comparison between these three extractors is made to analyze the goodness of EvoPruneDeepTL with them.

These features are used as the input for the fully-connected layers. We introduce two different compositions of these fully-connected layers:

- Single fully-connected layer: it is composed of a single layer with 512 neurons, followed by the output layer.
- Two fully-connected layers: this architecture has two layers of 512 each, and the output layer connecting the output of the last fully connected layer to as many neurons as the number of classes to be discriminated in the dataset.

EvoPruneDeepTL stands for the usage of sparse layers. By definition, a sparse layer has few active connections. A key object in this environment is the adjacency matrix. This matrix is key in our study because it is used to create a sparse layer from it. It allows decoding an individual evolved via the GA to yield, as a result, a neural network with a sparse layer. Taking a look about this matrix, EvoPruneDeepTL performs the pruning in relation to the connections/neurons that compose the input of the neuron. Consequently, there may be some neurons of the second layer which have no connection from the previous layer.

Based on the two network architectures described above, we consider different scenarios where EvoPruneDeepTL can be applied. We present these scenarios in the following figures, in which the red-dashed lines indicate the effects of the pruning.

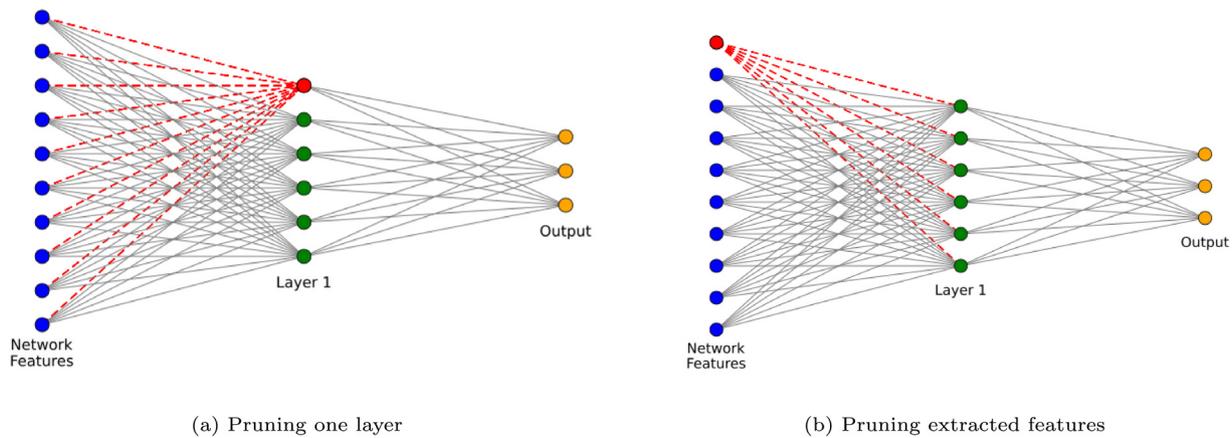


Fig. 4. Visualization of pruning architectures with one layer.

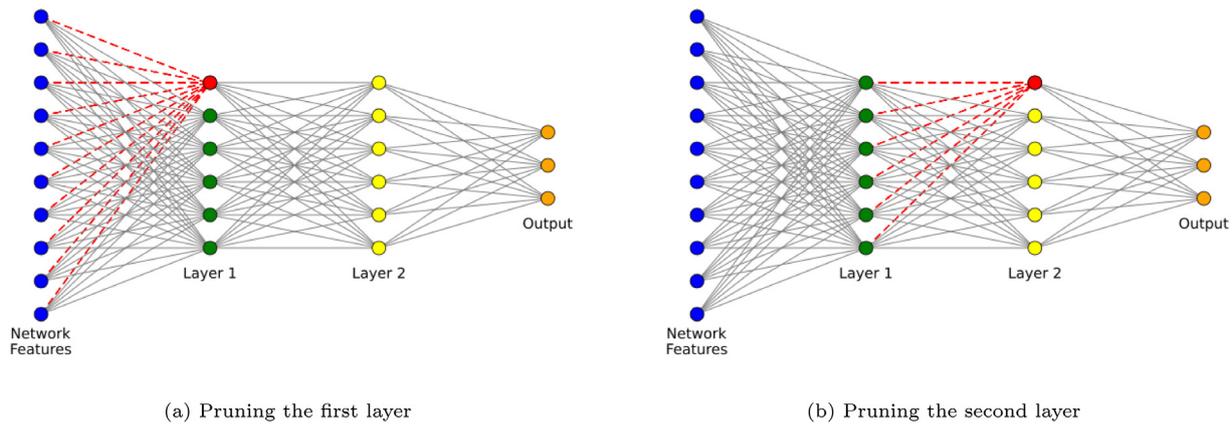


Fig. 5. Visualization of pruning architectures with two layers.

In addition to that, we have grouped the models in terms of the number of the last layers. The first model is the application of the EvoPruneDeepTL to prune model with one layer, which is shown in Fig. 4(a). Moreover, when the pruning is made with networks with two layers, three cases come up: pruning the first layer (see Fig. 5(a)), pruning the second layer (see Fig. 5(b) for these cases), or both at the same time, which is the combination of the last two cases. Lastly, EvoPruneDeepTL is also able to prune the characteristics that are extracted from the network. This approximation is called Feature Selection because EvoPruneDeepTL prunes the features that are less important to enhance the accuracy of the network. Fig. 4(b) illustrates how pruning in this last scenario reduces to a selection of features.

4. Experimental framework

In this section, we describe the experimental framework followed in our study. First, we give a brief description of the analyzed datasets. Then, the training setup is presented, emphasizing the parameters of EvoPruneDeepTL and the experimental conditions.

4.1. Datasets

In our study, we have chosen several diverse and representative datasets that are suitable for TL due to their size, as they require less training and inference time. Therefore, these datasets are suitable for population metaheuristics, as many individuals are evaluated. The selected datasets are shown in Table 1, which portrays their main characteristics for our experiments.

These datasets are diverse and taken from the literature:

Table 1

Datasets used in the experiments.

Dataset	Image size	# classes	# Instances (train/test)
SRSMAS	(299, 299)	14	333/76
RPS	(300, 300)	3	2520/372
LEAVES	(256, 256)	4	476/120
PAINTING	(256, 256)	5	7721/856
CATARACT	(256, 256)	4	480/121
PLANTS	(100, 100)	27	2340/236

- SRSMAS ([dataset] Gómez-Ríos et al., 2019) is a dataset to classify coral reef types with different classes and high distinction difficulty.
- RPS ([dataset] Laurence Moroney, 2019) is a dataset to identify the gesture of the hands in the popular Rock Paper Scissors game from images that have different positions and different skin colors.
- LEAVES is composed of images of healthy and unhealthy citrus leaves, with different shades of green ([dataset] Hafiz Tayyab Rauf et al., 2019).
- PLANTS is another dataset from the natural environment ([dataset] Singh et al., 2020, May), in which the task is to differentiate between leaves of different plants such as tomato, apple or corn, among others.
- CATARACT comes from the medical domain ([dataset] Sung-joon Choi, 2020), whose purpose is to classify different types of eye diseases.
- PAINTING is related to the painting world ([dataset] Virtual Russian Museum, 2018). The images in this dataset have



Fig. 6. Images of datasets. Top: SRSMAS examples. Middle: RPS examples. Bottom: LEAVES examples.

Table 2
Parameters of EvoPruneDeepTL.

Parameter	Value
Maximum Evals	200 (one layer) 300 (both layer)
# Runs	5
Population size	30
NAM	3
p_{mut}	0.07
Batch size	32

been taken from a museum and the task is to recognize different types of paintings.

Examples for several of the above datasets are shown in Fig. 6.

4.2. Training setup

The evaluation of EvoPruneDeepTL requires splitting the images of the datasets in train and test subsets. As the results could strongly depend on the train and test sets, we have applied in SRSMAS and LEAVES a 5-fold cross-validation²

For the remaining datasets, the train and test had already been defined beforehand, so we have used them for the sake of replicability.

The training is done using SGD as optimizer, a batch size of 32 images, and maximum 600 epochs, but the training stops when no improvement of loss is obtained in ten consecutive epochs. The model with the greater accuracy on the training set is saved. As we apply TL, only the last layers are trained, whereas the remaining ones are frozen with the parameter values imported from the pre-trained ResNet-50 network.

The parameters of EvoPruneDeepTL are indicated in Table 2. We have set the maximum evaluations to two different values,

² Sets for 5 fold CV for SRSMAS and LEAVES: https://drive.google.com/drive/folders/1Xf7OeZyWDDG-_Y4VX_nnAdfz3Kwhy8LU?usp=sharing. Last Access: 28/01/2022.

Table 3
Average time per run of EvoPruneDeepTL.

Dataset	First layer	Second layer	Both layers
SRSMAS	9 h 41 min	10 h 45 min	15 h 27 min
RPS	4 h 23 min	5 h 13 min	8 h 00 min
LEAVES	10 h 26 min	16 h 01 min	17 h 45 min
PAINTING	13 h 1 min	15 h 24 min	22 h 45 min
CATARACT	2 h 3 min	2 h 22 min	3 h 26 min
PLANTS	6 h 3 min	6 h 11 min	10 h 00 min

200 and 300, because there are some experiments we carry out to analyze the behavior of EvoPruneDeepTL that need an adaptation of this value because the search space in these experiments is wider. The size of the population of networks that our model evolves at each generation is set to 30, the mutation probability is p_{mut} and the NAM operator chooses the second parent among 3 candidates. The best solution found in terms of accuracy is returned. We note that in case of several solutions with the same accuracy, the returned solution is the configuration with the lowest percentage of active neurons. Note that the number of runs and total function evaluations is kept low to meet a computationally affordable balance between performance and the high execution times required for simulation. This is shown in Table 3, in which the average time per execution of the models with two layers is indicated. Unfortunately, this limited number of runs per experiment impedes the application of statistical tests to assess the significance of the reported differences, as tests conventionally used for this purpose require larger sample sizes to reach meaningful conclusions.

All the following experiments have been carried out using Python 3.6 and a Keras/Tensorflow implementation deployed and running on a Tesla V100-SXM2 GPU. The code is published in an open repository in GitHub.³

5. Results and discussion

In this section, we analyze the behavior of EvoPruneDeepTL. In order to show the benefits of using EvoPruneDeepTL, we propose four research questions (RQ) that they are going to be answered with different and diverse experiments over several datasets are carried out. We will show tables with the results of these experiments and we will analyze them to ensure the benefits of EvoPruneDeepTL. These RQ are the following ones:

(RQ1) Which is the performance of EvoPruneDeepTL against fully-connected models?

We compare EvoPruneDeepTL against non-pruned models comprising fully-connected layers to study which model obtains a better performance in the experiments. Moreover, we remark the flexibility of EvoPruneDeepTL applying it with one and several layers.

(RQ2) Which would be better, to remove neurons or connections?

We compare EvoPruneDeepTL using the two alternatives explained in the previous section: (1) pruning the neurons or (2) each individual represents exact connections between neurons, allowing for a more finely grained optimization. The goal of this section is to check which representation obtains the best results. On the one hand, the neuron representation of the length of chromosomes is shorter, so the domain search is smaller. On the other hand, the connections representation is a more fine-detail representation, so it could potentially allow the algorithm to obtain better results.

³ EvoPruneDeepTL repository: <https://github.com/ari-dasci/S-EvoDeepTLPruni> ng.

(RQ3) Which is the performance of EvoPruneDeepTL when compared to efficient pruning methods?

We compare the performance of EvoPruneDeepTL against several efficient pruning methods published in the literature for compressing CNN networks: Polynomial Decay (Zhu & Gupta, 2018, April), Weight Pruning (Han et al., 2015, December) and Neuron Pruning (Srinivas & Babu, 2015, September). This comparison of EvoPruneDeepTL and the CNN models is made in terms of accuracy and model compression.

(RQ4) Which would be better, pruning fully-connected layers or performing Feature Selection?

A particular case of EvoPruneDeepTL stands when the optimization of the network is done with one fully-connected layer and features are evolved towards the fittest for the problem at hand. Our aim is to check if this scheme improves the overall performance of the model over a dataset in terms of the accuracy of the models.

(RQ5) How does EvoPruneDeepTL perform when applied to different pre-trained networks?

We compare EvoPruneDeepTL with different feature extractors. From RQ1 to RQ4, ResNet-50 is used to extract the features or characteristics for the considered datasets. Experiments devised for this RQ aim to examine whether EvoPruneDeepTL adapts suitably to other feature extractors such as DenseNet-121 and VGG-19, so that better-performing pruned networks are produced by our proposal also for these feature extractors.

(RQ6) Can EvoPruneDeepTL adapt efficiently their pruned knowledge to changes in the modeling task, showing robustness?

We analyze the behavior of EvoPruneDeepTL when the datasets change. In this case, we have selected some datasets from our study, and we have done several modifications, removing partially or totally a class. Within these changes, we want to show both the robustness of the EvoPruneDeepTL in different situations and that the pruning models obtained by the GA of EvoPruneDeepTL have been adapted to each one of these situations.

This section is divided in Section 5.1, where the comparison of the diverse representations of pruning that EvoPruneDeepTL makes against the reference models is presented to answer RQ1. Next, Section 5.2 discusses whether EvoPruneDeepTL should operate over neurons or connections to analyze RQ2. Section 5.3 provides a complete comparison among EvoPruneDeepTL and other efficient pruning methods in order to solve RQ3. Section 5.4 explains the approximation of Feature Selection. A whole comparison against all the previous models is made to assess the importance of the Feature Selection to answer RQ4. Section 5.5 shows the comparison of the best two models of EvoPruneDeepTL with different feature extractors. Lastly, in Section 5.6, EvoPruneDeepTL is challenged, with several modifications of the used datasets, to grasp the relevant features of these datasets and to analyze the robustness of our proposal.

5.1. Answering RQ1: Pruning

In this section, we assess the performance gaps between the proposed EvoPruneDeepTL against other reference models to answer RQ1. In each subsection, several and diverse experiments are carried out to present results that assure the quality of EvoPruneDeepTL when it is compared to other models. This pruning section is composed of Section 5.1.1, in which we compare EvoPruneDeepTL against reference models with only one layer; of Section 5.1.2 we make the same experiments but with two layers

Table 4

Average results of EvoPruneDeepTL against not pruned models with one fully-connected layer.

Dataset	Measure	Not Pruned	Best Fixed	EvoPrune DeepTL
SRSMAS	Accuracy	0.832	0.866	0.885
	% Active neur.	100	20	25
RPS	Accuracy	0.938	0.938	0.954
	% Active neur.	100	40	46
LEAVES	Accuracy	0.923	0.927	0.935
	% Active neur.	100	80	38
PAINTING	Accuracy	0.939	0.945	0.951
	% Active neur.	100	60	46
CATARACT	Accuracy	0.703	0.719	0.732
	% Active neur.	100	70	39
PLANTS	Accuracy	0.432	0.432	0.480
	% Active neur.	100	10	49

and the optimization of one of them, and of Section 5.1.3, that shows the optimization of two consecutive layers at the same time.

In the following, we describe the different reference models:

- The first reference is composed of fully-connected layers of 512 units and the output layer. That is equivalent to the model with all neurons in active mode (all gens to 1). This model is the one with all active neurons, we call it *Not Pruned*.
- A grid search scheme is compared to EvoPruneDeepTL to check whether the improvement made by EvoPruneDeepTL could be obtained with a simple search over the percentage of neurons of the fully-connected layer. We have tested the fully-connected model with different number of neurons: 10% to 90% of its total units increasing this percentage by 10% (including both), and for each dataset we have identified the number of neurons which gives the best accuracy.
- The best result of the above models is also noted and it is called *Best Fixed*. When implemented over both layers, pruning is referred to as *Best Fixed Both*.

5.1.1. Pruning neurons of one fully-connected layer

This section introduces the results of pruning models with only one fully-connected layer. Table 4 shows the comparison of EvoPruneDeepTL against the reference models. In this case, the reference model only has one fully-connected layer composed by 512 units and the output layer.

For each dataset, the first row shows the obtained average accuracy by the models over the test set, whereas the second row informs about the average percentage of active neurons.

These results show how EvoPruneDeepTL is capable of distinguishing the pruning configurations that lead towards an improvement of performance of the models, as it obtains a greater accuracy in all the datasets for every reference model. Moreover, in most datasets, a higher compression ratio than the best fully-connected model is also achieved.

5.1.2. Pruning neurons of two fully-connected layers

In this section, our challenge is to improve the performance of a two fully-connected layer network. For that reason, EvoPruneDeepTL is applied to each layer individually.

The results of applying EvoPruneDeepTL to each layer individually are shown in Table 5, where **First Layer** indicates the case of the optimization of the first layer, and **Second Layer** describes the other case. In this case, Both *Not Pruned* and *Best Fixed* are the reference models with two fully-connected layer.

Table 5
Average results of EvoPruneDeepTL against not pruned models with two fully-connected layers.

Dataset	Measure	First layer			Second layer		
		Not Pruned	Best Fixed	EvoPrune DeepTL	Not Pruned	Best Fixed	EvoPrune DeepTL
SRSMAS	Accuracy	0.858	0.858	0.883	0.858	0.860	0.884
	% Active neur.	100	100	46	100	80	47
RPS	Accuracy	0.922	0.938	0.959	0.922	0.949	0.969
	% Active neur.	100	30	37	100	30	16
LEAVES	Accuracy	0.919	0.926	0.937	0.919	0.929	0.935
	% Active neur.	100	40	28	100	60	12
PAINTING	Accuracy	0.939	0.944	0.950	0.939	0.941	0.951
	% Active neur.	100	60	53	100	90	53
CATARACT	Accuracy	0.703	0.711	0.740	0.703	0.703	0.735
	% Active neur.	100	70	63	100	100	59
PLANTS	Accuracy	0.402	0.448	0.479	0.402	0.441	0.483
	% Active neur.	100	10	45	100	50	37

In this case, results follow the same path as the previous one: in all the datasets, EvoPruneDeepTL achieves an improvement of the accuracy over the reference models. Moreover, the Second Layer case obtains more compressed networks than the First Layer option.

Comparing the results of the scheme of one and two layers, both have similar results, only in RPS and CATARACT the difference in terms of accuracy is higher. Thus, these experiments have shown the ability of EvoPruneDeepTL of improving the overall performance of networks and, at the same time, reducing their complexity.

5.1.3. Pruning neurons of both layers

In the previous sections, we have tested EvoPruneDeepTL to single-layer optimization problems. In this section we increase the difficulty of the problem: the optimization of two consecutive fully-connected layers.

From the previous experiments, we have run EvoPruneDeepTL with 200 evaluations, but we have noticed that this number of evaluations might not be enough. This is due to the fact that we have now individuals of size 1024, 512 for each layer, and the search space is larger than in the rest of experiments. We have therefore also carried out the experiments with 300 function evaluations.

In Table 6, we show the results for reference models and EvoPruneDeepTL with 300 evaluations. The reference models stand the same as in the previous cases, but as they are implemented over both layers, the pruning is now referred to as *Best Fixed Both*.

In some cases, the percentage of remaining active neurons is higher than in the first and second layer models, but that is due to the complexity of this new problem. However, the performance of the network in these experiments indicates that the best option for pruning is achieved when the optimization is done to two consecutive layers.

These results prove the attainment of a sequential process to make pruning of DL models by adding layers and then, evolving their neurons to achieve a reduced configuration of the network. This process rises the performance of the models in terms of accuracy.

5.2. Answering RQ2: which would be better, to remove neurons or connections?

This section is devised to formally answer the RQ2, which is to decide if it is better to perform pruning of whole neurons or either single connections, by comparing different EvoPruneDeepTL chromosome representations: neurons and connections, as we described in Section 3. Two representations are shown in this

Table 6
Average results of EvoPruneDeepTL against not pruned methods optimizing two consecutive layers.

Dataset	Measure	Not Pruned	Best Fixed Both	EvoPrune DeepTL
SRSMAS	Accuracy	0.858	0.863	0.885
	% Active neur.	100	50	64
RPS	Accuracy	0.922	0.946	0.978
	% Active neur.	100	90	12
LEAVES	Accuracy	0.919	0.934	0.936
	% Active neur.	100	15	34
PAINTING	Accuracy	0.939	0.949	0.953
	% Active neur.	100	40	51
CATARACT	Accuracy	0.703	0.735	0.746
	% Active neur.	100	85	63
PLANTS	Accuracy	0.402	0.466	0.491
	% Active neur.	100	55	41

section: the neuron representation, in which a gen represents the connections of a neuron, and the connections representation, in which a gene represents a specific connection in the sparse layer. Neuron representation obtains shorter chromosomes than the connections one. Meanwhile, the connection representation leads to a more detailed representation and a larger domain search.

Table 7 shows for each dataset and representation the mean accuracy and % of active connections for both pruning methods. The connection strategy is named *Edges*. The results show that even though there are some cases in which the edges optimization achieves a similar performance of the network, the neuron optimization presents more robust results. The models working at the neuron level are even able to further reduce the number of active neurons in some datasets.

As a conclusion of this experiment, we can confirm that using the neuron approach is the best representation and that the second layer model gives us more consistent results than the first layer pruning model, both in accuracy and in reduction of the model.

5.3. Answering RQ3: Comparing EvoPruneDeepTL with efficient methods for CNN pruning

This section is devised to analyze the RQ3 comparing EvoPruneDeepTL to other well known network pruning methods to present results that measure the performance of our model against these methods. This comparison is conducted in terms of quality and computational complexity, aimed to prove the potential of EvoPruneDeepTL with respect to other pruning counterparts. To this end, we implement two different pruning methods,

Table 7
Average results of EvoPruneDeepTL against edges models.

Dataset	Measure	One layer		Two layers		
		Edges	EvoPruneDeepTL	Edges	EvoPruneDeepTL layer 1	EvoPruneDeepTL layer 2
SRSMAS	Accuracy	0.875	0.885	0.875	0.883	0.884
	% Active neur.	43	25	46	46	47
RPS	Accuracy	0.952	0.954	0.952	0.959	0.969
	% Active neur.	29	46	37	37	16
LEAVES	Accuracy	0.932	0.935	0.933	0.937	0.935
	% Active neur.	45	38	45	28	12
PAINTING	Accuracy	0.949	0.951	0.950	0.950	0.951
	% Active neur.	48	46	53	48	53
CATARACT	Accuracy	0.729	0.732	0.737	0.740	0.735
	% Active neur.	69	49	66	63	59
PLANTS	Accuracy	0.457	0.480	0.463	0.479	0.483
	% Active neur.	64	49	45	45	37

namely, weight pruning and neuron pruning. These methods have a parameter in common, $S_f \in \mathbb{R}(0, 1)$, which denotes the target pruning percentage. It is set to the same percentage of reduction that EvoPruneDeepTL has obtained in the experiments discussed previously. Next, we briefly describe each of such methods:

- **weight** (Han et al., 2015, December): Parameters with lower values are pruned at once. This method operates over the whole parameter set in the layer to be optimized. (Parameters: S_f)
- **polynomial decay** (Zhu & Gupta, 2018, April): Parameters are pruned guided by a Polynomial Decay schedule to the specified sparsity value. Between pruning steps, the network is allowed to fine tune for 5 epochs. This model is also applied over the whole parameter set in the layer to be optimized. Parameters used in the experimentation are listed in Table 8.
- **neuron** (Srinivas & Babu, 2015, September): Neurons with lower mean input connection values are pruned. (Parameters: S_f) as in Figs. 4 and 5. (Parameters: S_f)

Table 8 summarizes the value of the parameters of Polynomial Decay algorithm, which have been adapted to our experiments. Then, given a desired sparsity value of S , the sparsity is updated over a span of k pruning steps following the next equation

$$S_k = S_f + (S_i - S_f) \cdot \left(1 - \frac{K_k - K_i}{K_f - K_i}\right)^\alpha \text{ if } K_k \bmod F = 0 \quad (3)$$

wherein parameters are described as follows:

- $S_{i,f} \in \mathbb{R}(0, 1)$ are the initial and final sparsity percentages.
- S_f depends on the experiment. It is the percentage of pruning that EvoPruneDeepTL has achieved and the end of the generations.
- $K_{i,f} \in \mathbb{N}$ configures at what training step the pruning algorithm starts and ends.
- $K_k \in \mathbb{N}(K_i, K_f)$ is the current step.
- nb is the number of batches. It is calculated as the length of the training set divided by the batch size.
- F configures the frequency at which Eq. (3) is computed.

Parameters of the Polynomial Decay model are chosen to achieve a tradeoff between network recovery and the number of training epochs. Given the nature of this model, Polynomial Decay implies more training epochs than the implemented neuron and weight pruning methods. This fact could make the comparison between such methods unfair if the additional training epochs introduced by the Polynomial Decay model are high compared to the initial training epochs (i.e. 600). To avoid this situation,

Table 8
Parameter values of Polynomial Decay.

Parameter	Value
S_i	0.1
K_i	0
K_f	$nb \cdot 25$
F	$nb \cdot 5$
α	3.0

Table 9
Average results of EvoPruneDeepTL against efficient CNN for one layer models.

Dataset	Measure	One layer			
		Weight	Poly. Decay	Neuron	EvoPrune DeepTL
SRSMAS	Accuracy	0.805	0.823	0.745	0.885
	% Active neur.	25	25	25	25
RPS	Accuracy	0.917	0.927	0.869	0.954
	% Active neur.	46	46	46	46
LEAVES	Accuracy	0.918	0.920	0.886	0.935
	% Active neur.	38	38	38	38
PAINTING	Accuracy	0.993	0.994	0.874	0.951
	% Active neur.	46	46	46	46
CATARACT	Accuracy	0.678	0.679	0.658	0.732
	% Active neur.	39	39	39	39
PLANTS	Accuracy	0.406	0.411	0.365	0.480
	% Active neur.	49	49	49	49

Polynomial Decay is configured so that it sufficiently guarantees network recovery for all datasets while a minimal amount of extra training epochs are carried out, just an extra 4% from the initial 600 epochs (i.e. 25 extra epochs).

Our analysis aims to verify whether the performance of the above efficient pruning methods are comparable to EvoPruneDeepTL in terms of solution quality (accuracy) when they are configured to prune the same amount of parameters. Thus, the experimentation is carried out for the previously four cases discussed, selecting the average outcomes from the experimentation conducted in this point.

First, we show the results of this comparison when only a fully-connected layer is used for the optimization. Table 9 shows the results for this case. EvoPruneDeepTL outperforms the CNN models in five out of the six cases, but only in PAINTING these results are better for the Polynomial Decay or Weight models.

Second, Table 10 shows the results for the optimization of models with two layers, where **First Layer** indicates the cases of the optimization of the first layer and **Second Layer** describes the cases of the second layer. Results point out that EvoPruneDeepTL outperforms most of the methods in all the models and datasets.

Table 10
Average results of EvoPruneDeepTL against efficient CNN pruning methods for two layers models.

Dataset	Measure	First layer				Second layer				Both layers			
		Weight	Poly. Decay	Neuron	EvoPrune DeepTL	Weight	Poly. Decay	Neuron	EvoPrune DeepTL	Weight	Poly. Decay	Neuron	EvoPrune DeepTL
SRSMAS	Accuracy	0.795	0.815	0.775	0.883	0.834	0.837	0.779	0.884	0.845	0.847	0.647	0.885
	% Active neur.	46	46	46	46	47	47	47	47	64	64	64	64
RPS	Accuracy	0.886	0.911	0.803	0.959	0.845	0.911	0.696	0.969	0.694	0.899	0.490	0.978
	% Active neur.	37	37	37	37	16	16	16	16	12	12	12	12
LEAVES	Accuracy	0.913	0.918	0.812	0.937	0.904	0.919	0.712	0.935	0.911	0.925	0.747	0.936
	% Active neur.	28	28	28	28	12	12	12	12	34	34	34	34
PAINTING	Accuracy	0.995	0.993	0.850	0.950	0.937	0.938	0.920	0.951	0.934	0.940	0.853	0.953
	% Active neur.	53	53	53	53	53	53	53	53	51	51	51	51
CATARACT	Accuracy	0.668	0.684	0.673	0.740	0.694	0.689	0.648	0.737	0.686	0.696	0.611	0.746
	% Active neur.	63	63	63	63	59	59	59	59	63	63	63	63
PLANTS	Accuracy	0.408	0.403	0.343	0.479	0.392	0.420	0.313	0.482	0.393	0.411	0.278	0.491
	% Active neur.	45	45	45	45	37	37	37	37	41	41	41	41

Table 11
Times in seconds per run of EvoPruneDeepTL against efficient CNN pruning methods with one and two layers models.

Dataset	One layer				Two layers						
	Weight	Poly. Decay	Neuron	EvoPruneDeepTL	Weight	Poly. Decay	Neuron	EvoPruneDeepTL layer 1	EvoPruneDeepTL layer 2	EvoPruneDeepTL both layers	
SRSMAS	1995	2125	1995	34,510	2395	2545	2398	34,856	38,731	55,596	
RPS	1674	1893	1674	19,851	1229	1379	1229	15,758	18,790	28,774	
LEAVES	2425	2560	2425	35,243	2430	2565	2430	37,561	57,695	63,897	
PAINTING	1386	1508	1386	61,734	2903	3243	2903	46,856	55,414	81,913	
CATARACT	594	627	584	6,768	449	473	449	7,392	8,529	12,350	
PLANTS	298	407	298	28,456	270	370	270	21,788	22,235	35,998	

This case presents similar results as the one layer case because only in the PAINTING dataset EvoPruneDeepTL has a lower performance in relation to the literature methods. As a result of that, EvoPruneDeepTL’s robustness in performance over the literature methods has been shown in one-layer and two-layer networks.

Lastly, we compare the execution times for all the models. Evolutionary approaches are known to converge slowly in highly-dimensional search spaces, as the one tackled in this paper. For that reason, in this section, we also want to compare the required time of EvoPruneDeepTL and the other traditional approaches. Table 11 shows the time in seconds for each model. From these results, in terms of computational efficiency, our method suffers from the convergence slowness derived from the exploration of large search spaces.

To summarize, in this section we have fairly compared EvoPruneDeepTL to other well-known pruning methods, such as weight pruning and neuron pruning, guided by different pruning techniques. EvoPruneDeepTL is distinguished from other pruning methods due to the fact that they are advocate for shrinking the through their pruning process, but with an admissible decrease of the accuracy. Although our model is slower in terms of execution time, it scores higher accuracy levels than those of traditional pruning counterparts. Therefore, we conclude that EvoPruneDeepTL excels at determining which parameters to tune in neural networks with imported knowledge from other related tasks.

5.4. Answering RQ4: Feature Selection

The RQ4 establishes the dichotomy of choosing pruning or feature selection for the given problem. For that reason, in this section we analyze the FS model by conducting the same group of experiments of the previous sections, to compare it against EvoPruneDeepTL to decide which one scores best among them. The FS scheme is a particular case of EvoPruneDeepTL if only one fully-connected layer composes the configuration of the network

Table 12
Average results for Feature Selection against non pruning methods.

Dataset	Measure	Not Pruned	Best Fixed	Feature Selection
SRSMAS	Accuracy	0.832	0.866	0.884
	% Active neur.	100	20	60
RPS	Accuracy	0.938	0.938	0.985
	% Active neur.	100	40	45
LEAVES	Accuracy	0.923	0.927	0.943
	% Active neur.	100	80	59
PAINTING	Accuracy	0.939	0.945	0.958
	% Active neur.	100	60	55
CATARACT	Accuracy	0.703	0.719	0.747
	% Active neur.	100	70	55
PLANTS	Accuracy	0.432	0.432	0.472
	% Active neur.	100	10	68

and the pruning and GA are focused on the extracted features of the ResNet-50 model.

Table 12 shows the results for this model against the reference methods. This case follows the same similarities of the previous ones, as FS obtains the best average results for all the datasets.

Similarly to the previous sections, we have also compared this model with the CNN pruning methods with only one layer, as shown in Table 13. In this case, in four out of six datasets the Feature Selection outperforms these methods, but in LEAVES and PAINTING Weight and Polynomial Decay perform better than our model.

In this section, we have compared our FS scheme against reference models and efficient pruning methods published in the literature. The results shed light over the benefits of this model as it is also able to achieve a great performance over the reference models and also, in most cases, against the CNN pruning methods.

The global results for EvoPruneDeepTL and its different versions are presented in Table 14. The rows show the achieved

Table 13
Average results of feature selection against efficient CNN pruning methods.

Dataset	Measure	Feature Selection			
		Weight	Poly. Decay	Neuron	Feature Selection
SRSMAS	Accuracy	0.841	0.878	0.802	0.884
	% Active neur.	60	60	60	60
RPS	Accuracy	0.913	0.926	0.869	0.985
	% Active neur.	45	45	45	45
LEAVES	Accuracy	0.947	0.940	0.946	0.943
	% Active neur.	59	59	59	59
PAINTING	Accuracy	0.962	0.968	0.883	0.958
	% Active neur.	55	55	55	55
CATARACT	Accuracy	0.696	0.689	0.687	0.747
	% Active neur.	55	55	55	55
PLANTS	Accuracy	0.421	0.317	0.402	0.472
	% Active neur.	68	68	68	68

Table 14
Results and percentage of improvement for each version of EvoPruneDeepTL in relation to each reference model.

Dataset	Measure	Pruning model one layer	Pruning model both layers	Feature Selection
SRSMAS	Accuracy	0.885	0.885	0.884
	% Improvement	1.9	2.2	1.8
RPS	Accuracy	0.954	0.978	0.985
	% Improvement	1.6	3.2	4.7
LEAVES	Accuracy	0.935	0.936	0.943
	% Improvement	0.8	0.2	1.6
PAINTING	Accuracy	0.951	0.953	0.958
	% Improvement	0.6	0.4	1.3
CATARACT	Accuracy	0.732	0.746	0.747
	% Improvement	1.3	1.1	2.8
PLANTS	Accuracy	0.480	0.491	0.472
	% Improvement	4.8	2.5	4.0

accuracy and the percentage of improvement in relation to the best reference models for each model.

Reviewing the results of EvoPruneDeepTL and FS, we confirm that FS is the best model, as it obtains the best accuracy levels in four out of six datasets. Furthermore, the pruning of both layers carried out by EvoPruneDeepTL also attains very notable performance levels.

Moreover, if we consider the optimization using the pruning model, the optimization of both layers yields the best results in terms of mean accuracy for each dataset. However, when comparing pruning and FS, the latter has more robust models: it achieves the best performance in four datasets, and it is also shown in the improvement percentage for each dataset.

In conclusion, it is shown with empirical evidence that pruning can be done by optimizing the fully-connected layers, specifically, by evolving their neurons to get the fittest configuration that reports an improvement of the network performance. An evolutionary feature selection based on the extracted features also achieves a great network performance, both in improving the accuracy and in reducing its complexity.

5.5. Answering RQ5: Comparing different EvoPruneDeepTL with different feature extractors

This section is devised to formally compare EvoPruneDeepTL with different networks that serve as feature extractor for each dataset to analyze. CNNs have shown their capability to overcome different and diverse classification problems by learning visual features that best correlate with the target variable of the task

at hand. Transferring this knowledge to other problems with a similar domain, which is what TL stands for, also helps to the capability of generalization of the model that is devised for the target task, specially when the volume of data for that task is low.

In previous sections, we have observed that the combination of EvoPruneDeepTL with ResNet-50 has improved both reference models and pruning methods of the literature. However, in this section, we explore two other feature extractors and assess whether such performance gaps prevail. To determine performance gains of EvoPruneDeepTL when using these alternative feature extractors, both are tested over the two best performing scenarios of EvoPruneDeepTL, namely pruning both layers and feature selection. Moreover, we also include in the comparison these feature extractors with pruning algorithms from the literature, following the same experimental procedure described in preceding sections.

The chosen feature extractors are DenseNet-121 and VGG-19. The experiments with these networks have been carried out in the same conditions that the previous ones have been done. Table 15 shows the comparison of these networks against the reference models based on fully-connected layers.

The results from the previous table show the ability of EvoPruneDeepTL to adapt itself to several feature extractors. For both DenseNet-121 and VGG-19, it improves the reference models. Taking a deep look at the three networks, ResNet-50 is the best of them, as it has the best improvement over several datasets. Nonetheless, the straight conclusion which is derived from these experiments is that EvoPruneDeepTL is able to adapt to different feature extractors and datasets.

Once we have seen that EvoPruneDeepTL has achieved better results than the reference models, now we inspect the performance of different pruning methods from the literature. For that reason, we compare the model which prunes two consecutive layers against its similar models from the pruning methods. Table 16 shows the comparison of the three networks against these methods. Note that, for this comparison, we have used the best to models of the CNN pruning methods: Weights and Polynomial Decay. This holds for the rest of this section.

The results shown in Table 16 give rise to interesting insights. To begin with, the first three columns are related to ResNet-50, which have a great performance over these reference models, and we know it from the previous sections. However, DenseNet and VGG are totally new in this kind of experiments. The reality is that both of these networks improve the CNN pruning methods when they are applied to fully-connected layers in most cases. Only in a few experiments are better than obtained by our proposal.

A global vision of these experiments suggests that DenseNet and VGG, just like ResNet, contribute to the discovery of pruned neural networks that maximize accuracy and reduce the number of active neurons. Moreover, these results verify that EvoPruneDeepTL is able to achieve for different networks better results than reference and efficient CNN pruning methods when the pruning is made in two consecutive layers. Thus, we have shown the adaption ability of EvoPruneDeepTL for this case using several networks (ResNet, DenseNet, and VGG).

The following scenario is the feature selection model which derives from EvoPruneDeepTL. This scenario, which encourages the pruning of the features extracted by the pre-trained network, has yielded the best results so far. Next, Table 17 shows the comparison of these networks when feature selection is performed.

In the previous sections, we have shown that EvoPruneDeepTL is able to prune the extracted features derived from the network, and this model has reached the best results of EvoPruneDeepTL. Table 17 shows that the pruning of the features that they have been extracted using different networks (DenseNet and VGG) also increases the performance of the networks, which is shown in

Table 15
Average results of EvoPruneDeepTL with different networks optimizing two consecutive layers against non pruning methods.

Dataset	Measure	ResNet-50			DenseNet-121			VGG19		
		Not Pruned	Best Fixed Both	EvoPrune DeepTL	Not Pruned	Best Fixed Both	EvoPrune DeepTL	Not Pruned	Best Fixed Both	EvoPrune DeepTL
SRSMAS	Accuracy	0.858	0.863	0.885	0.861	0.881	0.890	0.837	0.853	0.885
	% Active neur.	100	50	64	100	50	72	100	85	55
RPS	Accuracy	0.922	0.946	0.978	0.704	0.723	0.754	0.814	0.879	0.922
	% Active neur.	100	90	12	100	70	43	100	40	50
LEAVES	Accuracy	0.919	0.934	0.936	0.896	0.904	0.915	0.903	0.911	0.917
	% Active neur.	100	15	34	100	60	39	100	50	68
PAINTING	Accuracy	0.939	0.949	0.953	0.940	0.943	0.947	0.923	0.938	0.945
	% Active neur.	100	40	51	100	40	67	100	70	34
CATARACT	Accuracy	0.703	0.735	0.746	0.694	0.727	0.741	0.661	0.727	0.759
	% Active neur.	100	85	63	100	55	51	100	45	42
PLANTS	Accuracy	0.402	0.466	0.491	0.411	0.428	0.456	0.292	0.364	0.374
	% Active neur.	100	55	41	100	55	78	100	50	64

Table 16
Average results of EvoPruneDeepTL against efficient CNN pruning methods for pruning consecutive layers.

Dataset	Measure	ResNet50			DenseNet-121			VGG19		
		Weight	Poly. Decay	EvoPrune DeepTL	Weight	Poly. Decay	EvoPrune DeepTL	Weight	Poly. Decay	EvoPrune DeepTL
SRSMAS	Accuracy	0.845	0.847	0.885	0.862	0.865	0.890	0.933	0.869	0.885
	% Active neur.	64	64	64	72	72	72	55	55	55
RPS	Accuracy	0.694	0.899	0.978	0.815	0.817	0.754	0.803	0.842	0.922
	% Active neur.	12	12	12	43	43	43	50	50	50
LEAVES	Accuracy	0.911	0.925	0.936	1.000	0.907	0.915	1.000	0.916	0.917
	% Active neur.	34	34	34	39	39	39	68	68	68
PAINTING	Accuracy	0.934	0.940	0.953	0.897	0.901	0.947	0.923	0.922	0.945
	% Active neur.	51	51	51	67	67	67	34	34	34
CATARACT	Accuracy	0.686	0.696	0.746	0.587	0.593	0.741	0.661	0.686	0.759
	% Active neur.	63	63	63	51	51	51	42	42	42
PLANTS	Accuracy	0.393	0.411	0.491	0.251	0.249	0.456	0.284	0.292	0.374
	% Active neur.	41	41	41	78	78	78	64	64	64

Table 17
Average results for feature selection with different networks against non-pruning methods.

Dataset	Measure	ResNet-50			DenseNet-121			VGG19		
		Not Pruned	Best Fixed	Feature Selection	Not Pruned	Best Fixed	Feature Selection	Not Pruned	Best Fixed	Feature Selection
SRSMAS	Accuracy	0.832	0.866	0.884	0.858	0.881	0.896	0.753	0.766	0.869
	% Active neur.	100	20	60	100	60	68	100	40	87
RPS	Accuracy	0.938	0.938	0.985	0.720	0.720	0.839	0.887	0.890	0.982
	% Active neur.	100	40	45	100	100	48	100	60	53
LEAVES	Accuracy	0.923	0.927	0.943	0.896	0.902	0.921	0.852	0.876	0.924
	% Active neur.	100	80	59	100	10	60	100	10	68
PAINTING	Accuracy	0.939	0.945	0.958	0.934	0.941	0.956	0.924	0.924	0.943
	% Active neur.	100	60	55	100	90	63	100	100	77
CATARACT	Accuracy	0.703	0.719	0.747	0.669	0.702	0.787	0.628	0.661	0.765
	% Active neur.	100	70	55	100	30	57	100	10	66
PLANTS	Accuracy	0.432	0.432	0.472	0.394	0.415	0.464	0.335	0.352	0.376
	% Active neur.	100	10	68	100	90	67	100	40	66

the accuracy of these networks over the datasets. There are some cases in which non pruning methods have less active neurons, but their accuracy is lower than the models of EvoPruneDeepTL. For that reason, the feature selection keeps being the best of EvoPruneDeepTL models, as it has the best results so far.

The next, and final, step is to check the performance of this feature selection model against efficient CNN pruning methods from the literature. In this section, we have checked that, for models which prune two consecutive layers, EvoPruneDeepTL performs better than the pruning methods. Consequently, now we focus on this comparison, but in terms of the models which

prune the extracted features. [Table 18](#) shows the results of this comparison for the different networks.

The results show, not only that ResNet-50 has a great performance (same results as previous sections), that both DenseNet and VGG outperform the pruning methods when applied to prune the features extracted from the networks. Both new networks show a better performance in all the datasets than the reference methods. For that reason, we conclude that the usage of EvoPruneDeepTL with these three networks has proven the capability to perform better than the pruning methods.

In this section, we have compared ResNet-50 with other two networks in two different scenarios: pruning consecutive layers

Table 18
Average results of EvoPruneDeepTL with different networks against efficient CNN pruning methods for feature selection models.

Dataset	Measure	ResNet50			DenseNet-121			VGG19		
		Weight	Poly. Decay	Feature Selection	Weight	Poly. Decay	Feature Selection	Weight	Poly. Decay	Feature Selection
SRSMAS	Accuracy	0.841	0.878	0.884	0.869	0.868	0.896	0.826	0.825	0.869
	% Active neur.	60	60	60	68	68	68	87	87	87
RPS	Accuracy	0.913	0.926	0.985	0.675	0.699	0.839	0.981	0.834	0.982
	% Active neur.	45	45	45	48	48	48	53	53	53
LEAVES	Accuracy	0.947	0.940	0.943	0.858	0.891	0.921	0.904	0.843	0.924
	% Active neur.	59	59	59	60	60	60	68	68	68
PAINTING	Accuracy	0.962	0.968	0.958	0.937	0.934	0.956	0.928	0.928	0.943
	% Active neur.	55	55	55	63	63	63	77	77	77
CATARACT	Accuracy	0.696	0.689	0.747	0.676	0.682	0.787	0.666	0.688	0.765
	% Active neur.	55	55	55	57	57	57	66	66	66
PLANTS	Accuracy	0.421	0.317	0.472	0.387	0.394	0.464	0.322	0.311	0.376
	% Active neur.	68	68	68	67	67	67	66	66	66

and pruning the extracted features from the networks. The experiments show that EvoPruneDeepTL has proven its ability to adapt to other networks in both cases and has improved both reference models and pruning methods of the literature. For that reason, and in light of the results from the previous sections, we conclude that EvoPruneDeepTL has shown the ability of facing diverse tasks, as EvoPruneDeepTL has achieved a great performance when different networks are used either for pruning consecutive layer or pruning the features extracted from the network.

5.6. Answering RQ6: Analyzing the ability of EvoPruneDeepTL to adapt to relevant classes and robustness

The purpose of this section is twofold. First, we want to analyze the goodness of EvoPruneDeepTL when modeling varying problems. In this case, we want to see how it adapts to the different classes that make up the datasets so that it captures the relevance of each of them. For a given dataset, we analyze each of its classes to determine if EvoPruneDeepTL is also able to have a good performance over it, and then, compare these results with the whole dataset.

The second objective of this section is the analysis of these results. Once EvoPruneDeepTL has modeled each of the classes for a dataset at hand, we check the quality of the obtained results. For that reason, we must check that results are not affected by the stochasticity induced by the usage of a genetic algorithm at the core of the proposed EvoPruneDeepTL. Recall that stochasticity implies that the output of the algorithm may not be the same, even with the same input. For that reason, our second objective is related to this factor, and we want to show that the effect of randomness has a low impact on EvoPruneDeepTL, i.e., the good results do not depend on the randomness.

In order to measure the effect of randomness in the pruned networks evolved by EvoPruneDeepTL, we resort to a similarity measure called Centered Kernel Alignment, CKA (Kornblith et al., 2019, June). CKA measures the similarity of trained neural networks, in compliance with several invariance properties that must be met for these particular computing structures (namely invariance to invertible linear transformation, invariance to orthogonal transformation and invariance to isotropic scaling). We compare the trained networks as a result of the application of EvoPruneDeepTL. This comparison answers the question about the robustness of EvoPruneDeepTL.

CKA is based on the Hilbert–Schmidt Independence Criterion, HSIC (Gretton et al., 2005, October). It compares two matrices (**K** and **L**) and determines the level of independence between them, as it is shown in (4). In our case, these matrices are the structures that contain the weights of the trained neural networks. CKA

takes a maximum value of 1 when the two inputs of CKA are the same matrix. The range of this measure is [0, 1]. This means that both matrix are very similar (in that case because they are identical). Thus, if the CKA value is high, then both matrix are similar.

$$CKA(K, L) = \frac{HSIC(K, L)}{\sqrt{HSIC(K, K)HSIC(L, L)}} \tag{4}$$

CKA is a measure which allows us make a double comparison. Note that this measure helps us to compare the genotype (chromosomes of the GA) against the phenotype (pruned networks of EvoPruneDeepTL). The interpretation of this measure follows that if the similarity of the chromosomes is high, then the CKA value should be also high (close to one, which is its maximum). Five independent executions have been made of EvoPruneDeepTL, so we have taken the output of each of them, and we have performed a double comparison based on this trained neural network, which is explained next:

1. Comparison against the closest element ($CKA_{Closest}$): we have the *Output* element as the best one for an execution. Then, we calculate the Hamming distance of the best element with respect to all the elements evaluated in the evolutionary process of EvoPruneDeepTL. The element with the smallest Hamming distance to the best one is denoted as *Closest*. Then, *Output* is compared against *Closest*. The robustness of EvoPruneDeepTL is tested in this comparison because high values of CKA when similar chromosomes are compared is essential, as it will show that the results are not due to randomness, but to the process that EvoPruneDeepTL performs.
2. Comparison against fully-connected, reference models (CKA_{Ref}): in this case, we compare *Output* against a fully-connected network with all the neurons activated. This comparison sheds light on the ability of EvoPruneDeepTL to learn which neurons are the best to solve the problem at hand. It also permits to explain the difference in terms of accuracy between the models that EvoPruneDeepTL develops against the reference models. This value of CKA quantifies the differences in accuracy between our models and the reference models, as EvoPruneDeepTL searches for the best neurons to remove the unnecessary ones, while reference models simply train the models without taking into account the neurons which should be removed.

In this section, we select the best two models of EvoPruneDeepTL for these experiments: pruning consecutive layers (both case) and pruning the features extracted from the

Table 19
Average results of CATARACT-Class with pruning consecutive layers and feature selection.

Dataset	Measure	EvoPruneDeepTL both	No pruning both	EvoPruneDeepTL feature selection	No pruning feature selection
CATARACT - Retina	Accuracy	0.844	0.801	0.871	0.697
	% Active neur.	43	100	54	100
CATARACT - Glaucoma	Accuracy	0.846	0.789	0.857	0.554
	% Active neur.	44	100	63	100
CATARACT - Cataract	Accuracy	0.735	0.732	0.761	0.614
	% Active neur.	76	100	49	100
CATARACT - Normal	Accuracy	0.833	0.805	0.843	0.655
	% Active neur.	53	100	57	100
CATARACT - Full	Accuracy	0.746	0.703	0.747	0.703
	% Active neur.	63	100	55	100

networks (feature selection). Moreover, three datasets are considered in the experiments designed for this section: CATARACT, PAINTING and RPS. Lastly, we show different tables with the following structure: *DATASET-Class*. This means that the mentioned *DATASET* is analyzed without the class called *Class*.

For the CKA comparison, we will show two groups of tables, one per each type of model (pruning consecutive layer and feature selection). Moreover, each table is composed of the problem at hand and the mean values of the Hamming distance and CKA averaged over five executions of EvoPruneDeepTL. The rows of each table correspond to each *DATASET-Class* and the four columns represent in pairs the previously explained comparisons, $CKA_{Closest}$ and CKA_{Ref} , as we show the mean Hamming distance and its corresponding CKA mean value for each of the comparisons.

5.6.1. Analyzing the relevance of each class for a given dataset

Given a dataset of n classes, this approach performs a procedure that removes a whole class of the dataset and then, EvoPruneDeepTL is applied with that remaining data both in pruning consecutive layers and in feature selection. As a result of that process, n experiments are done for each dataset.

The structure of the these tables correspond with the usual structure of the rest of the paper, but now we show four columns. These columns represent, in pairs based on the model, the results of EvoPruneDeepTL versus the reference model without pruning. The difference between the pairs of columns is the model at hand: pruning consecutive layers or feature selection. We note that for the first two columns, the optimization process which lies in EvoPruneDeepTL is made with 300 evaluations and two-layers networks, meanwhile for the last columns, the process of pruning the extracted features is made with only one-layer networks and 200 evaluations, i.e., under the same conditions as the experiments in the previous sections.

The first results show the CATARACT dataset under these conditions. Table 19 shows these results. We conclude from these results that Cataract class is the easiest class in the dataset, as both models struggle with that class (third row of the table), but they improve their results with it. However, the results of EvoPruneDeepTL are better than the reference model. The number of active neurons at the end of the evolutionary process are reduced, in most cases, nearly by half. Moreover, the feature selection model is also able to decrease this number from its previous results (55% of remaining active neurons) in some cases. The same conclusion can be drawn in relation to pruning consecutive layers, as the full dataset has a mean percentage of active neurons of 63% and in three of four cases this number is reduced. The results show that the pruning of the extracted features of the network, i.e., the feature selection which derives from EvoPruneDeepTL, is the best approximation for this dataset.

We focus now on the following dataset, RPS. Table 20 shows the results of these experiments. Results show that the *Paper*

class makes an easier dataset, as all the approaches reach the maximum accuracy. The other two experiments show that EvoPruneDeepTL achieves a better performance to the reference models. For pruning consecutive layers, the number of remaining active neurons is higher in comparison with the full dataset, but in the feature selection model this number is very similar or even for RPS-Scissors is lower.

The last considered dataset is PAINTING. Table 21 shows the results for each of the experiments which have been made for this dataset. The table shows that both of EvoPruneDeepTL models are constantly achieving better results than the reference models. However, the difference between the accuracy is higher in the feature selection models than the pruning consecutive layers. Taking into consideration the remaining active neurons, the feature selection has a similar degree of pruning in relation with the experiments which have been carried out for the full dataset experiments (55%). Similar conclusion can be drawn for the other model, as the mean percentage of active neuron for pruning consecutive layers is 51% and we have models with fewer active neurons, but also with higher percentage.

These experiments shed light on a conclusion that it is not far from the previous sections. The feature selection model, which performs the pruning of the extracted features of the network, constitutes the best model for all the experiments, as it has the most difference between EvoPruneDeepTL and reference models. Nonetheless, the effect of pruning consecutive layers is also positive, as it is shown in the results of the previous tables.

The next part of this section is crucial to determine the robustness of EvoPruneDeepTL. Moreover, the differences in accuracy of the previous tables are going to be explained in the following tables. The key element of the comparison is the CKA measure and the Hamming distance of the solutions. The combination of these values determine the key points that we have discussed at the beginning of this section.

Therefore, we are showing the CKA tables for each experiment to perform the commented double comparison in this section. These tables have a different structure from the last tables, so we explain how to interpret them. Both of them present a similar structure, but the table which shows the results for pruning consecutive layers has another column. This column shows the value of the CKA measure for the layer at hand. In the case of the prune of the extracted feature of the network (feature selection), as they only has one layer, this column is not required.

We show this pair of tables for each *DATASET-Class*. The first table shows the results of the feature selection model, and the second one presents the results of the pruning of consecutive layers.

The composition of the tables for feature selection models is the following one. After the first column which show the dataset at hand, the next two columns represent the comparison among the phenotype and genotype of EvoPruneDeepTL, i.e, chromosomes and pruned networks. Both mean Hamming distance of

Table 20
Average results of general RPS pruning consecutive layers and feature selection.

Dataset	Measure	EvoPruneDeepTL both	No pruning both	EvoPruneDeepTL feature selection	No pruning feature selection
RPS - Paper	Accuracy	1.000	1.000	1.000	1.000
	% Active neur.	51	100	51	100
RPS - Rock	Accuracy	0.985	0.979	1.000	0.996
	% Active neur.	51	100	49	100
RPS - Scissors	Accuracy	0.994	0.955	1.000	0.955
	% Active neur.	33	100	23	100
RPS Full	Accuracy	0.978	0.922	0.985	0.938
	% Active neur.	12	100	45	100

Table 21
Average results of general PAINTING pruning consecutive layers and feature selection.

Dataset	Measure	EvoPruneDeepTL both	No pruning both	EvoPruneDeepTL feature selection	No pruning feature selection
PAINTING - Sculpture	Accuracy	0.942	0.932	0.947	0.857
	% Active neur.	64	100	53	100
PAINTING - Painting	Accuracy	0.959	0.946	0.961	0.820
	% Active neur.	42	100	50	100
PAINTING - Iconography	Accuracy	0.942	0.920	0.949	0.835
	% Active neur.	30	100	57	100
PAINTING - Engraving	Accuracy	0.979	0.974	0.979	0.883
	% Active neur.	74	100	56	100
PAINTING - Drawings	Accuracy	0.994	0.989	0.996	0.944
	% Active neur.	57	100	57	100
PAINTING Full	Accuracy	0.953	0.939	0.958	0.939
	% Active neur.	51	100	55	100

the five executions and $CKA_{Closest}$ are shown. The following group of two columns shows the other explained comparison of the reference models, which have all the neurons active. The metrics are the same as the previous case, but now the CKA (CKA_{Ref}) corresponds to the mean value from the best model to this reference model.

The composition of the table for the pruning of consecutive layers is similar to the previous case. However, another column is required for a more detailed explanation. This column gives information about the CKA value for the layer at hand. Due to the fact that we are comparing the whole chromosome, both Hamming distance values are common to both layers, but the CKA value is layer dependent. Then, in this table, we want to highlight that the best chromosome EvoPruneDeepTL obtains good values of the CKA measure for each of the layers of the model.

First, we show the CKA values for the CATARACT dataset in its four different cases of DATASET-Class. Tables 22 and 23 show the results for feature selection and pruning consecutive layers, respectively. The results show for both models the robustness of EvoPruneDeepTL because the mean CKA of the best versus its closest chromosome in the history is a value extremely close to 1, which is the maximum value (this value is reached when the best solution is compared to itself).

Taking a more deep look at the results of the feature selection, we see that the Hamming distance is very low, which is a good result and also proves the robustness of EvoPruneDeepTL. The second group of columns, in which the comparison is made against a model with all neurons active, we see that the Hamming distance is higher and this has an impact on the CKA value, which is higher. The conclusion which derives from these experiments is that EvoPruneDeepTL learns to distinguish the valuable neurons which have an impact on the model. This CKA value is the explanation of the difference in accuracy in the previous experiments of this dataset between EvoPruneDeepTL models and reference models.

The results that they are shown in Table 23 confirm the robustness of EvoPruneDeepTL. This insight is the same as in

Table 22
Comparison of the CKA measure for feature selection in CATARACT-Class.

Dataset	Feature Selection			
	Hamming distance EvoPruneDeepTL	$CKA_{Closest}$	Hamming distance no pruned model	CKA_{Ref}
CATARACT - Retina	0.005	0.981	0.457	0.283
CATARACT - Glaucoma	0.002	0.991	0.372	0.376
CATARACT - Cataract	0.001	0.994	0.514	0.200
CATARACT - Normal	0.011	0.961	0.423	0.250

the previous table: low values of Hamming distance in the best chromosomes of EvoPruneDeepTL implies high values of CKA in the closest element. Moreover, the Hamming distance from the elements of pruning consecutive layers is higher than in the other case, but the CKA values are also higher. This is a fair result because the difference in accuracy between EvoPruneDeepTL and the reference models is lower than in the other case. Note that the class Cataract from this dataset is the class with the fewest gap in accuracy, and this is shown in its CKA value. The Glaucoma class is the opposite of Cataract, and the CKA value is lower. In all the cases, EvoPruneDeepTL confirms the ability to learn the neurons that they are indispensable to achieve a greater performance, and that is the main difference between the reference models.

The second dataset under analysis is RPS. Table 24 shows the results of the feature selection models for RPS-Class. We see that the pair Hamming distance and CKA of the closest have a great result in two of the three cases. Moreover, the results of the other metrics achieve a great results, similarly to CATARACT-Class with this model. For that reason, we confirm that, for this model, EvoPruneDeepTL is also able to learn the neurons that maximize the accuracy for the problem.

Table 23
Comparison of the CKA measure for pruning consecutive layers in CATARACT-Class.

Dataset	Pruning consecutive layers				
	# Layer	Hamming distance EvoPruneDeepTL	CKA _{closest}	Hamming distance no pruned model	CKA _{Ref}
CATARACT - Retina	Layer 1	0.020	0.980	0.567	0.685
	Layer 2		0.983		0.769
CATARACT - Glaucoma	Layer 1	0.001	0.996	0.556	0.680
	Layer 2		0.997		0.745
CATARACT - Cataract	Layer 1	0.001	0.998	0.240	0.891
	Layer 2		0.998		0.916
CATARACT - Normal	Layer 1	0.027	0.975	0.469	0.784
	Layer 2		0.979		0.843

Table 24
Comparison of the CKA measure for feature selection in RPS-Class.

Dataset	Feature Selection			
	Hamming distance EvoPruneDeepTL	CKA _{closest}	Hamming distance no pruned model	CKA _{Ref}
RPS - Paper	0.199	0.416	0.485	0.224
RPS - Rock	0.020	0.926	0.507	0.195
RPS - Scissors	0.040	0.863	0.766	0.104

A special case is RPS-Paper. When RPS does not have this class, the problem seems to be a very easy task, because all the models in the previous experiments for this dataset achieve the maximum accuracy. This is the only case in which the CKA for the best and its closest element is lower in comparison with the others. This is due to the fact that the problem at hand can be solved with many chromosomes, as they all have the maximum accuracy value, so the chromosomes might not be very similar, because the range of possible solutions is wide.

Table 25 shows the pruning of consecutive layers that it is performed by EvoPruneDeepTL. Similar conclusions are obtained from these results. First, we see that the Paper class has the same problem which appears in the previous table. However, the other two groups of experiments are harder to solve and this have been drawn in the CKA of the closest element, because both layers have a great value of this measure. In the counterpart, the CKA values for the reference also has a similar understanding, which belongs to the fact that EvoPruneDeepTL pruning of consecutive layers learns the best neurons for both layers.

In overall, EvoPruneDeepTL is also a robust model for this dataset both in feature selection and pruning consecutive layers, and it also proves that the difference in accuracy between our models and the reference models is stated in the CKA_{Ref}.

The last dataset in this section is PAINTING. The first table relates to the feature selection model of EvoPruneDeepTL. Table 26, once more, shows that the lowest Hamming distance of the best chromosome when it is compared with its closest, brings high values of CKA. The conclusion is clear, EvoPruneDeepTL is robust. Moreover, the values of CKA_{Ref} are also a good estimation of how EvoPruneDeepTL looks for the best neurons. Both CATARACT and PAINTING have lots of similarities in the feature selection model.

Table 27 shows the results for pruning consecutive layers in the dataset PAINTING. These results, again, prove that the closest and best elements of EvoPruneDeepTL achieve a great value of CKA given a low value of Hamming distance, which is the best output that we can have. Moreover, the CKA values for the reference models are high, but this is due to the fact that the difference in accuracy between the models is lower. However, this also proves that EvoPruneDeepTL also learns the best neurons for this problem.

In this section, we have shown that EvoPruneDeepTL is able to capture the relevance of the diverse classes and datasets that they are shown. Thanks to that adaption, EvoPruneDeepTL has shown its robustness and its ability to search for the best neurons to tackle the problem at hand.

5.6.2. Effects of a gradual aggregation of a class in the problem at hand

This section is devised to analyze the impact of a class when it appears as a new class in a dataset, and it increases its number of examples over time. We have selected the class Iconography of PAINTING for these experiments. The reason lies in the fact that the class with a low percentage of examples is a minority class of this dataset, but it becomes the majority class of PAINTING when all the examples are used. Adding more examples of this class lets us check how EvoPruneDeepTL is able to adapt to different scenarios for the same dataset when a class is gradually growing on its importance in the dataset.

This section has a similar structure to the previous one. First, we show the results of the experiments for pruning consecutive layers and feature selection models of EvoPruneDeepTL. Then, the CKA values is also presented to perform the same double comparison as it has been done in the last section. For this section, the notation for the dataset is PAINTING-Iconography we talk about the dataset resulting from adding more examples, and in the tables this is shown as PAINTING-Pct%, where Pct = 20, 40, 60 and 80.

The first experiments we show are in Table 28. These experiments have been carried out with the dataset PAINTING with the different percentage, as we have previously explained. The results show that, as the percentage of data increases, the models tend to become better. If we compare these results with those obtained with the full dataset, we see that the model with 80% of the data (and with 60% of the data also for feature selection) is the closest to the full model (see Tables 6 and 12).

Reviewing the results tables with the full dataset and comparing them with these results, we observe that for the consecutive layer pruning model, the number of active neurons is lower in these experiments. The same phenomenon occurs in most of the feature selection cases, except when 40% of the data is used, where this number increases as the model improves the accuracy for that dataset at the cost of increasing the percentage of active neurons.

Next, we follow the same process as for the previous section, in which the different classes of various datasets were analyzed. We show the value of CKA for the feature selection model and for the pruning consecutive layers model. Table 29 shows the results for the different feature selection models applied to the various data percentage options of the Iconography class. The CKA value, which compares the best with its closest element in the history of the execution, is very high. This implies that EvoPruneDeepTL is a robust model, since the phenotype obtained from the genotype

Table 25
Comparison of the CKA measure for pruning consecutive layers in RPS-Class.

Dataset	Pruning consecutive layers				
	# Layer	Hamming distance EvoPruneDeepTL	CKA _{Closest}	Hamming distance no pruned model	CKA _{Ref}
RPS - Paper	Layer 1	0.217	0.725	0.486	0.764
	Layer 2		0.761		0.825
RPS - Rock	Layer 1	0.032	0.977	0.492	0.757
	Layer 2		0.983		0.852
RPS - Scissors	Layer 1	0.010	0.982	0.667	0.597
	Layer 2		0.974		0.668

Table 26
Comparison of the CKA measure for feature selection in PAINTING-Class.

Dataset	Feature Selection			
	Hamming distance EvoPruneDeepTL	CKA _{Closest}	Hamming distance no pruned model	CKA _{Ref}
PAINTING - Sculpture	0.002	0.992	0.471	0.259
PAINTING - Painting	0.003	0.986	0.502	0.202
PAINTING - Iconography	0.002	0.993	0.433	0.236
PAINTING - Engraving	0.006	0.984	0.441	0.291
PAINTING - Drawings	0.021	0.925	0.432	0.236

is very similar. In addition, the other CKA value reported by the reference model indicates that EvoPruneDeepTL is capable of the neurons important for the model, thus explaining the difference in accuracy between the two approaches.

The following table, Table 30. The robustness of the proposal becomes evident when comparing the best element with its closest element in each of the runs, which has also occurred in the previous case. CKA values are extremely high when these elements are compared.

The comparison with respect to the reference models shows results similar to those of other cases of consecutive layer pruning. The difference in accuracy is reflected in the CKA, which is higher than in the feature selection cases, because this difference is larger when it comes to the pruning of the extracted features of the network.

This section has allowed us to see EvoPruneDeepTL in different situations it has had to face. From datasets with fewer classes so that our proposal is able to adapt to all the subclasses that compose it (first subsection), to the gradual increase of a class from being the minority to the majority (second subsection). The study which has been performed in this section relies on a measure which allows us to study the robustness of EvoPruneDeepTL and, in addition, allows us to see the differences in accuracy of the models that they have been developed.

The results in both sets of experiments show that the stochasticity that might be present in the proposal is not influential. The results of the CKA measure when comparing the best trained network found by EvoPruneDeepTL and its closest trained network for each of the runs show the high degree of robustness of EvoPruneDeepTL.

The comparison of EvoPruneDeepTL with networks with all neurons active (CKA_{Ref}) shows us a twofold conclusion. When we are dealing with the feature selection models, those that performed the pruning of the extracted features of the network, the difference in accuracy is reflected in the value of CKA, which is very low and that means that the models are very different. On the other hand, for the case of pruning consecutive layers, the

CKA_{Ref} value reflects models with less difference in comparison to the previous case. However, both in feature selection and pruning consecutive layers it is observed that EvoPruneDeepTL is able to search for the neurons that best approximate the problem to be solved.

6. Advantages and disadvantages of EvoPruneDeepTL

This section is devised to discuss the advantages and disadvantages of EvoPruneDeepTL, considering the diverse and large experimentation which has been done in the manuscript. The advantages of EvoPruneDeepTL can be summarized as follows:

- **Specialization of the last layers of networks.**

An important element of EvoPruneDeepTL is the transfer learning. This is one of the most commonly used techniques. We have refined its process, which is the extraction of pre-trained features and then, the specialization of the fully-connected layers. In this context, EvoPruneDeepTL, and specifically the GA which is composed of, when applied to these layers does not limit the network learning compared to other evolutionary models in the literature that require high computational time to evaluate the datasets, as they train the whole network. For that reason, EvoPruneDeepTL can be applied to more complex datasets.

- **Performance over reference models and efficient pruning methods from the literature.**

The usage of an evolutionary model that focuses on pruning neurons of the fully-connected layers achieves a better performance than other pruning methods when applied under the same conditions of EvoPruneDeepTL. The positive effect of the genetic algorithm is the selection of the best neurons of these layers, so that the evolution towards the best configuration for the networks is obtained thanks to EvoPruneDeepTL.

- **Constructive modeling over the last layers of the networks.**

In the different experiments that have been carried out in the sections, we have observed that performing the pruning constructively based on the number of layers achieves good results. Pruning one-layer networks achieves good results, but when the number of layer increases, it is shown that performing the pruning over a single layer of two-layer networks improves the one-layer networks. Nonetheless, the simultaneous pruning of the both layers achieves a better modeling of the datasets than all the previous pruning models.

- **Pruning the extracted features of the network against pruning fully-connected layers.**

EvoPruneDeepTL obtains better results by pruning the self-generated features resulting from transfer learning versus pruning the fully-connected layers. This is an intuitive idea because the learned patterns or features are different for each problem, and the learned features for the original problem may not be useful for the target problem. Knowing

Table 27
Comparison of the CKA measure for pruning consecutive layers in PAINTING-Class.

Dataset	Pruning consecutive layers				
	# Layer	Hamming distance EvoPruneDeepTL	CKA _{closest}	Hamming distance no pruned model	CKA _{Ref}
PAINTING - Sculpture	Layer 1	0.001	0.998	0.359	0.832
	Layer 2		0.999		0.879
PAINTING - Painting	Layer 1	0.004	0.998	0.578	0.701
	Layer 2		0.998		0.777
PAINTING - Iconography	Layer 1	0.001	0.998	0.701	0.593
	Layer 2		0.998		0.672
PAINTING - Engraving	Layer 1	0.001	0.999	0.255	0.895
	Layer 2		0.999		0.931
PAINTING - Drawings	Layer 1	0.003	0.998	0.427	0.797
	Layer 2		0.998		0.874

Table 28
Average results for PAINTING-Iconography with pruning consecutive layers and feature selection.

Dataset	Measure	EvoPruneDeepTL both	No pruning both	EvoPruneDeepTL feature selection	No pruning feature selection
PAINTING - 20%	Accuracy	0.944	0.931	0.945	0.826
	% Active neur.	47	100	56	100
PAINTING - 40%	Accuracy	0.945	0.929	0.950	0.826
	% Active neur.	39	100	65	100
PAINTING - 60%	Accuracy	0.946	0.931	0.954	0.839
	% Active neur.	53	100	51	100
PAINTING - 80%	Accuracy	0.947	0.927	0.953	0.839
	% Active neur.	41	100	47	100
PAINTING Full	Accuracy	0.953	0.939	0.958	0.939
	% Active neur.	51	100	55	100

Table 29
Comparison of the CKA measure for feature selection in PAINTING-Iconography.

Dataset	Feature Selection			
	Hamming distance EvoPruneDeepTL	CKA _{closest}	Hamming distance no pruned model	CKA _{Ref}
PAINTING - 20%	0.001	0.995	0.442	0.277
PAINTING - 40%	0.007	0.976	0.352	0.461
PAINTING - 60%	0.019	0.934	0.484	0.205
PAINTING - 80%	0.001	0.998	0.530	0.177

which characteristics matter is crucial to the problem at hand. The evolutionary process allows pruning these features and selecting those that best solve the modeling problem under consideration.

• **Generalization of EvoPruneDeepTL to other feature extractors.**

One of the advantages of EvoPruneDeepTL is that the model is generalizable to diverse feature extractors. This is an important advantage because EvoPruneDeepTL is able to achieve a great performance over different datasets and with diverse networks. These are used to extract the features of the dataset, thanks to the transfer learning technique. EvoPruneDeepTL has improved the reference models and pruning models from the literature in both the pruning of two consecutive layers and the pruning of the features extracted by the networks.

• **Adaptation to relevant classes and gradual aggregation of data for the problem at hand.**

EvoPruneDeepTL has achieved a great performance in different situations. However, it is also important to see how

it adapts to different situations within the datasets themselves. EvoPruneDeepTL has improved the performance of the reference models, but a measure is needed to support the quality of these models. However, this increase in performance has been proven by CKA not to be the result of chance, so EvoPruneDeepTL models are able to search for the best neurons to maximize the accuracy of the problem.

• **The good results of EvoPruneDeepTL do not depend on randomness.**

The evolutionary algorithm in which EvoPruneDeepTL lies in is a stochastic algorithm and the results may be affected by randomness. This is a risk about the model, so it is required to check if the good results are biased by the randomness. For that reason, we have introduced a new measure, CKA. This measure compares the differences between the pruned networks, and it is a way to measure the robustness of EvoPruneDeepTL. The values of the CKA for the various experiments of the previous section shed light on the fact that EvoPruneDeepTL results do not depend on the randomness.

The main disadvantage of EvoPruneDeepTL is:

• **Execution time of EvoPruneDeepTL.**

The main drawback of EvoPruneDeepTL it is the time that is required to execute the model. In comparison with the pruning methods from the literature and the reference models, the table of execution times (see Table 11) shows the speed of the other models, but EvoPruneDeepTL is slower than these models. The time difference is made up by improved network performance, thanks to the usage of EvoPruneDeepTL. Nonetheless, there are some practical cases in which the training time is not a problem, like in medical diagnosis, because the main objective is obtaining a better percentage of the models in terms of accuracy.

Table 30
Comparison of the CKA measure for pruning consecutive layers in PAINTING-Iconography.

Dataset	Pruning consecutive layers				
	# Layer	Hamming distance EvoPruneDeepTL	CKA _{Closest}	Hamming distance no pruned model	CKA _{Ref}
PAINTING - 20%	Layer 1	0.001	0.999	0.529	0.732
	Layer 2		0.998		0.794
PAINTING - 40%	Layer 1	0.001	0.998	0.607	0.677
	Layer 2		0.999		0.755
PAINTING - 60%	Layer 1	0.001	0.999	0.466	0.758
	Layer 2		0.999		0.828
PAINTING - 80%	Layer 1	0.001	0.998	0.594	0.694
	Layer 2		0.998		0.762

7. Conclusions

This paper has introduced EvoPruneDeepTL, a novel model that sparsifies the architecture of the last layers of a DL model initialized using TL. EvoPruneDeepTL is a combination of sparse layers and EA, so that the neurons of these layers are pruned using the EA, in order to adapt them to the problem to tackle and deciding which neurons/connections to leave active or inactive.

EvoPruneDeepTL is a flexible model that optimizes models with one and two layers and even two layers at the same time. Our results show that the pruning over complete neurons is better than pruning connections individually, establishing the last one as the best encoding strategy. The evolution of the sparse layer improves these models in terms of accuracy and also in terms of complexity of the network. In comparison with compared reference models and pruning methods from the literature, EvoPruneDeepTL achieves a better performance than all of them. The choice of one among the pruning models or feature selection has been answered and informed with experimental evidence: the FS scheme derived from EvoPruneDeepTL has shown a better performance, in most cases, than the pruning methods. The ability of adaptation of EvoPruneDeepTL to other feature extractors has been tested. Lastly, EvoPruneDeepTL has also shown its capability to adapt to the relevance of diverse problems and it has also achieved an outstanding level of robustness, which implies that the results do not depend on random nature of the search operators used by the GA that lies at the core of the proposed evolutionary pruning method.

From an overarching perspective, this work aligns with a growing strand of contributions where evolutionary computation and DL have synergized together to yield optimized models that attain better levels of performance and/or an increased computational efficiency. Indeed, this fusion of concepts (forged as Evolutionary Deep Learning) has been used for other optimization processes, including hyperparameter or structural tuning. Another recent case of the symbiosis of EA's and DL are represented in AutoML-Zero that use an evolutionary search to automatically search the best DL structure. AutoML-Zero and EvoPruneDeepTL are two great examples of the benefits of combining EA's and DL that outline the potential and promising path of successes envisioned for this research area.

Future research work stemming from the results reported in this study is planned from a two-fold perspective. To begin with, we plan to achieve larger gains from the combination of DL and EA by extending the evolutionary search over higher layers of the neural hierarchy, increasing the number of optimized layers and neurons per layer. To this end, we envision that exploiting the layered arrangement in which neurons are deployed along the neural architecture will be essential to ensure an efficient search. The second research line relates to this last thought, aiming to improve the search algorithm itself by resorting to advanced concepts in evolutionary computation (e.g. niching methods or co-evolutionary algorithms).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Both data and code are freely available online and they are referenced in the manuscript.

Acknowledgments

F. Herrera, D. Molina and J. Poyatos are supported by the Andalusian Excellence project P18-FR-4961, the infrastructure project with reference EQC2018-005084-P and the R&D and Innovation project with reference PID2020-119478GB-I00 granted by the Spain's Ministry of Science and Innovation and European Regional Development Fund (ERDF). Aritz D. Martinez and Javier Del Ser would like to thank the Basque Government for the funding support received through the EMAITEK and ELKARTEK programs, as well as the Consolidated Research Group MATHMODE (IT1456-22) granted by the Department of Education of this institution, Spain.

References

- Aneja, N., & Aneja, S. (2019, July). Transfer learning using CNN for handwritten devanagari character recognition. In *1st international conference on advances in information technology (ICAIT)*, Chikmagalur, India, 2019.
- Anwar, S., Hwang, K., & Sung, W. (2017). Structured pruning of deep convolutional neural networks. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 13(3), Article 32. <http://dx.doi.org/10.1145/3005348>.
- Assunção, F., Lourenço, N., Machado, P., & Ribeiro, B. (2019). DENSER: deep evolutionary network structured representation. *Genetic Programming and Evolvable Machines*, 20(1), 5–35. <http://dx.doi.org/10.1007/s10710-018-9339-y>.
- Back, T., Fogel, D. B., & Michalewicz, Z. (1997). *Handbook of evolutionary computation* (1st ed.). IOP Publishing Ltd..
- Chambers, L. D. (2000). *The practical handbook of genetic algorithms: Applications* (2nd ed.). Chapman and Hall/CRC.
- [dataset] Gómez-Ríos, A., Tabik, S., Luengo, J., Shihavuddin, A., & Herrera, F. (2019). Coral species identification with texture or structure images using a two-level classifier based on convolutional neural networks. *Knowledge-Based Systems*, 184, Article 104891. <http://dx.doi.org/10.1016/j.knsys.2019.104891>.
- [dataset] Hafiz Tayyab Rauf, Saleem, B. A., Lali, M. I. U., Khan, M. A., Sharif, M., & Bukhari, S. A. C. (2019). A citrus fruits and leaves dataset for detection and classification of citrus diseases through machine learning. *Data in Brief*, 26, Article 104340. <http://dx.doi.org/10.1016/j.dib.2019.104340>.
- [dataset] Laurence Moroney (2019). Rock, paper, scissors dataset. Retrieved from <http://www.laurencemoroney.com/rock-paper-scissors-dataset/>. Accessed September 10, 2020.
- [dataset] Singh, D., Jain, N., Jain, P., Kayal, P., Kumawat, S., & Batra, N. (2020, May). Plantdoc: a dataset for visual plant disease detection. In *7th ACM IKDD CoDS and 25th COMAD, Hyderabad, India, 2020*.
- [dataset] Sungjoon Choi (2020). Cataract dataset. Retrieved from <https://www.kaggle.com/jr2ngb/cataractdataset>. Accessed September 10, 2020.

- [dataset] Virtual Russian Museum (2018). Art images: drawing/painting/sculptures/engravings. Retrieved from <https://www.kaggle.com/thedownhill/art-images-drawings-painting-sculpture-engraving>, Accessed September 10, 2020.
- Elsken, T., Metzger, J. H., & Hutter, F. (2019a, May). Efficient multi-objective neural architecture search via Lamarckian evolution. In *7th international conference on learning representations (ICLR)*, New Orleans, LA, USA, 2019.
- Elsken, T., Metzger, J. H., & Hutter, F. (2019b). Neural architecture search: A survey. *Journal of Machine Learning Research*, 20(1), 1997–2017, <https://dl.acm.org/doi/10.5555/3322706.3361996>.
- Fernandes, C., & Rosa, A. (2001, May). A study on non-random mating and varying population size in genetic algorithms using a royal road function. In *Proceedings of the 2001 congress on evolutionary computation, Seoul, Korea, 2001*.
- Frankle, J., & Carbin, M. (2019, May). The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *7th international conference on learning representations (ICLR)*, New Orleans, LA, USA, 2019.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization and machine learning* (1st ed.). Addison-Wesley Longman Publishing Co., Inc.
- Gómez-Ríos, A., Tabik, S., Luengo, J., Shihavuddin, A., Krawczyk, B., & Herrera, F. (2019). Towards highly accurate coral texture images classification using deep convolutional neural networks and data augmentation. *Expert Systems with Applications*, 118, 315–328. <http://dx.doi.org/10.1016/j.eswa.2018.10.010>.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning* (1st ed.). MIT Press.
- Gretton, A., Bousquet, O., Smola, A., & Schölkopf, B. (2005, October). Measuring statistical dependence with Hilbert-Schmidt norms. In *International conference on algorithmic learning theory, Singapore, 2005*.
- Han, S., Mao, H., & Dally, W. (2016, May). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *4th international conference on learning representations (ICLR)*, San Juan, Puerto Rico, 2016.
- Han, S., Pool, J., Tran, J., & Dally, W. (2015, December). Learning both weights and connections for efficient neural network. In *Proceedings of the 28th international conference on neural information processing systems, Montreal Canada, 2015*.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016, June). Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition, Las Vegas, NV, USA, 2016*.
- Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017, July). Densely connected convolutional networks. In *IEEE conference on computer vision and pattern recognition (CVPR)*, Honolulu, HI, USA, 2017.
- Iba, H. (2018). *Evolutionary approach to machine learning and deep neural networks: Neuro-evolution and gene regulatory networks* (1st ed.). Springer.
- Iguyon, J., & Elisseff, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3, 1157–1182, <https://dl.acm.org/doi/10.5555/944919.944968>.
- Jung, S., Park, J., & Lee, S. (2019, May). Polyphonic sound event detection using convolutional bidirectional LSTM and synthetic data-based transfer learning. In *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, Brighton, UK, 2019.
- Khan, S., Islam, N., Jan, Z., Din, I. U., & Rodrigues, J. J. C. (2019). A novel deep learning based framework for the detection and classification of breast cancer using transfer learning. *Pattern Recognition Letters*, 125, 1–6. <http://dx.doi.org/10.1016/j.patrec.2019.03.022>.
- Kokiopoulou, E., Hauth, A., Sbaiz, L., Gesmundo, A., Bartók, G., & Berent, J. (2020, August). Task-aware performance prediction for efficient architecture search. In *24th European conference on artificial intelligence, Santiago de Compostela, Spain, 2020*.
- Kornblith, S., Norouzi, M., Lee, H., & Hinton, G. (2019, June). Similarity of neural network representations revisited. In *Proceedings of the 36th international conference on machine learning, California, CA, USA, 2019*.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012, December). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems, Lake Tahoe, NV, USA, 2012*.
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324. <http://dx.doi.org/10.1109/5.726791>.
- Lee, H., Pham, P., Largman, Y., & Ng, A. Y. (2009, December). Unsupervised feature learning for audio classification using convolutional deep belief networks. In *Advances in neural information processing systems, Vancouver, B.C., Canada, 2009*.
- Liu, H., Simonyan, K., Vinyals, O., Fernando, C., & Kavukcuoglu, K. (2018, April). Hierarchical representations for efficient architecture search. In *6th international conference on learning representations (ICLR)*, Vancouver, B.C., Canada, 2018.
- Liu, Z., Sun, M., Zhou, T., Huang, G., & Darrell, T. (2019, May). Rethinking the value of network pruning. In *7th international conference on learning representations (ICLR)*, New Orleans, LA, USA, 2019.
- Liu, W., Wang, Z., Liu, X., Zeng, N., Liu, Y., & Alsaadi, F. (2017). A survey of deep neural network architectures and their applications. *Neurocomputing*, 234, 11–26. <http://dx.doi.org/10.1016/j.neucom.2016.12.038>.
- Liu, J., Wang, Y., & Qiao, Y. (2017, February). Sparse deep transfer learning for convolutional neural network. In *31st AAAI conference on artificial intelligence (AAAI)*, San Francisco, CA, USA.
- Liu, C., Zoph, B., Shlens, J., Hua, W., Li, L.-J., Fei-Fei, L., Yuille, A., Huang, J., & Murphy, K. (2018, September). Progressive neural architecture search. In *European conference on computer vision, Munich, Germany 2018*.
- Long, X., Ben, Z., & Liu, Y. (2019). A survey of related research on compression and acceleration of deep neural networks. *Journal of Physics: Conference Series*, 1213, Article 052003. <http://dx.doi.org/10.1088/1742-6596/1213/5/052003>.
- Long, X., Ben, Z., Zeng, X., Liu, Y., Zhang, M., & Zhou, D. (2019). Learning sparse convolutional neural network via quantization with low rank regularization. *IEEE Access*, 7, 51866–51876. <http://dx.doi.org/10.1109/ACCESS.2019.2911536>.
- Lu, Z., Sreekumar, G., Goodman, E., Banzhaf, W., Deb, K., & Boddeti, V. (2021). Neural architecture transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43, 2971–2989. <http://dx.doi.org/10.1109/TPAMI.2021.3052758>.
- Lu, Z., Whalen, I., Boddeti, V., Dhebar, Y., Deb, K., Goodman, E., & Banzhaf, W. (2019, July). NSGA-Net: Neural architecture search using multi-objective genetic algorithm. In *GECCO 2019 - Proceedings of the 2019 genetic and evolutionary computation conference, Prague, Czech Republic, 2019*.
- Luo, J.-H., Wu, J., & Lin, W. (2017, October). ThiNet: A Filter Level Pruning Method for Deep Neural Network Compression. In *IEEE international conference on computer vision, Venice, Italy, 2017*.
- Mantzaris, D., Anastassopoulos, G., & Adamopoulos, A. (2011). Genetic algorithm pruning of probabilistic neural networks in medical disease estimation. *Neural Networks*, 24, 831–835. <http://dx.doi.org/10.1016/j.neunet.2011.06.003>.
- Martín, A., Lara-Cabrera, R., Fuentes-Hurtado, F., Naranjo, V., & Camacho, D. (2018). EvoDeep: a new evolutionary approach for automatic deep neural networks parametrisation. *Journal of Parallel and Distributed Computing*, 117, 180–191. <http://dx.doi.org/10.1016/j.jpdc.2017.09.006>.
- Martinez, A. D., Del Ser, J., Villar-Rodríguez, E., Osaba, E., Poyatos, J., Tabik, S., Molina, D., & Herrera, F. (2021). Lights and shadows in Evolutionary Deep Learning: Taxonomy, critical methodological analysis, cases of study, learned lessons, recommendations and challenges. *Information Fusion*, 67, 161–194. <http://dx.doi.org/10.1016/j.inffus.2020.10.014>.
- Masson, H., Bhuian, A., Nguyen-Meidine, L. T., Javan, M., Siva, P., Ayed, I. B., & Branger, E. (2021). Exploiting prunability for person re-identification. *EURASIP Journal on Image and Video Processing*, 2021, Article 22. <http://dx.doi.org/10.1186/s13640-021-00562-6>.
- Mehdipour Ghazi, M., Yanikoglu, B., & Aptoula, E. (2017). Plant identification using deep neural networks via optimization of transfer learning parameters. *Neurocomputing*, 235, 228–235. <http://dx.doi.org/10.1016/j.neucom.2017.01.018>.
- Mohapatra, R., Saha, S., Coello, C. A. C., Bhattacharya, A., Dhavala, S. S., & Saha, S. (2022). AdaSwarm: Augmenting gradient-based optimizers in deep learning with swarm intelligence. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 6, 329–340. <http://dx.doi.org/10.1109/TETCI.2021.3083428>.
- Muhammad, K., Khan, S., Ser, J. D., & Albuquerque, V. H. C. d. (2021). Deep learning for multigrade brain tumor classification in smart healthcare systems: A prospective survey. *IEEE Transactions on Neural Networks and Learning Systems*, 32, 507–522. <http://dx.doi.org/10.1109/TNNLS.2020.2995800>.
- Muhammad, K., Ullah, A., Lloret, J., Ser, J. D., & de Albuquerque, V. H. C. (2020). Deep learning for safe autonomous driving: Current challenges and future directions. *IEEE Transactions on Intelligent Transportation Systems*, 22(7), 4316–4336. <http://dx.doi.org/10.1109/TITS.2020.3032227>.
- Nibali, A., He, Z., & Wollersheim, D. (2017). Pulmonary nodule classification with deep residual networks. *International Journal of Computer Assisted Radiology and Surgery*, 12(10), 1799–1808. <http://dx.doi.org/10.1007/s11548-017-1605-6>.
- Noda, K., Yamaguchi, Y., Nakadai, K., Okuno, H. G., & Ogata, T. (2015). Audio-visual speech recognition using deep learning. *Applied Intelligence*, 42(4), 722–737. <http://dx.doi.org/10.1007/s10489-014-0629-7>.
- Pan, S., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359. <http://dx.doi.org/10.1109/TKDE.2009.191>.
- Real, E., Aggarwal, A., Huang, Y., & Le, Q. V. (2019, January). Regularized evolution for image classifier architecture search. In *Proceedings of the AAAI conference on artificial intelligence, Honolulu, HI, USA, 2019*.
- Real, E., Liang, C., So, D., & Le, Q. (2020, July). Automl-zero: evolving machine learning algorithms from scratch. In *International conference on machine learning, 2020*.
- Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y. L., Tan, J., Le, Q. V., & Kurakin, A. (2017, August). Large-scale evolution of image classifiers. In *International conference on machine learning, Sydney, Australia, 2017*.

- Roy, D., Murty, K. S. R., & Mohan, C. K. (2015, July). Feature selection using deep neural networks. In *2015 international joint conference on neural networks (IJCNN), Killarney, Ireland, 2015*.
- Sa, I., Ge, Z., Dayoub, F., Upcroft, B., Perez, T., & McCool, C. (2016). Deepfruits: A fruit detection system using deep neural networks. *Sensors*, 16(8), Article 1222. <http://dx.doi.org/10.3390/s16081222>.
- Salehinejad, H., & Valaee, S. (2021). Edropout: Energy-based dropout and pruning of deep neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 1–14. <http://dx.doi.org/10.1109/TNNLS.2021.3069970>, in press.
- Samala, R., Chan, H.-P., Hadjiiski, L., Helvie, M., Richter, C., & Cha, K. (2018). Evolutionary pruning of transfer learned deep convolutional neural network for breast cancer diagnosis in digital breast tomosynthesis. *Physics in Medicine and Biology*, 63(9), Article 095005. <http://dx.doi.org/10.1088/1361-6560/aabb5b>.
- Scott, G. J., England, M. R., Starns, W. A., Marcum, R. A., & Davis, C. H. (2017). Training deep convolutional neural networks for land-cover classification of high-resolution imagery. *IEEE Geoscience and Remote Sensing Letters*, 14(4), 549–553. <http://dx.doi.org/10.1109/LGRS.2017.2657778>.
- Shin, H.-C., Roth, H., Gao, M., Lu, L., Xu, Z., Nogue, I., Yao, J., Mollura, D., & Summers, R. (2016). Deep Convolutional Neural Networks for Computer-Aided Detection: CNN Architectures, Dataset Characteristics and Transfer Learning. *IEEE Transactions on Medical Imaging*, 35(5), 1285–1298. <http://dx.doi.org/10.1109/TMI.2016.2528162>.
- Simonyan, K., & Zisserman, A. (2015, May). Very deep convolutional networks for large-scale image recognition. In *3rd international conference on learning representations, San Diego, CA, USA, 2015*.
- Srinivas, S., & Babu, R. V. (2015, September). Data-free parameter pruning for deep neural networks. In *Proceedings of the British machine vision conference (BMVC), Swansea, UK, 2015*.
- Stanley, K., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2), 99–127. <http://dx.doi.org/10.1162/106365602320169811>.
- Sultana, F., Sufian, A., & Dutta, P. (2018, November). Advancements in image classification using convolutional neural network. In *2018 fourth international conference on research in computational intelligence and communication networks (ICRCICN), Kolkata, India, 2018*.
- Tajbakhsh, N., Shin, J. Y., Gurudu, S. R., Hurst, R. T., Kendall, C. B., Gotway, M. B., & Liang, J. (2016). Convolutional neural networks for medical image analysis: Full training or fine tuning? *IEEE Transactions on Medical Imaging*, 35(5), 1299–1312. <http://dx.doi.org/10.1109/TMI.2016.2535302>.
- Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., & Liu, C. (2018, October). A survey on deep transfer learning. In *International conference on artificial neural networks (ICANN), Rhodes, Greece, 2018*.
- Ullrich, K., Welling, M., & Meeds, E. (2017, April). Soft weight-sharing for neural network compression. In *5th international conference on learning representations (ICLR), Toulon, France, 2017*.
- Wan, S., Liang, Y., & Zhang, Y. (2018). Deep convolutional neural networks for diabetic retinopathy detection by image classification. *Computers & Electrical Engineering*, 72, 274–282. <http://dx.doi.org/10.1016/j.compeleceng.2018.07.042>.
- Wang, Z., Li, F., Shi, G., Xie, X., & Wang, F. (2020). Network pruning using sparse learning and genetic algorithm. *Neurocomputing*, 404, 247–256. <http://dx.doi.org/10.1016/j.neucom.2020.03.082>.
- Weiss, K., Khoshgoftaar, T. M., & Wang, D. (2016). A survey of transfer learning. *Journal of Big Data*, 3(1), Article 9. <http://dx.doi.org/10.1186/s40537-016-0043-6>.
- Wen, L., Li, X., Li, X., & Gao, L. (2019, May). A new transfer learning based on VGG-19 network for fault diagnosis. In *IEEE 23rd international conference on computer supported cooperative work in design (CSCWD), Porto, Portugal, 2019*.
- Yildirim, Ö., Pławiak, P., Tan, R.-S., & Acharya, U. (2018). Arrhythmia detection using deep convolutional neural network with long duration ECG signals. *Computers in Biology and Medicine*, 102, 411–420. <http://dx.doi.org/10.1016/j.combiomed.2018.09.009>.
- Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014, December). How transferable are features in deep neural networks? In *Advances in neural information processing systems, Montreal, Canada, 2014*.
- Zhou, X., Gong, W., Fu, W., & Du, F. (2017, May). Application of deep learning in object detection. In *IEEE/ACIS 16th international conference on computer and information science (ICIS), Wuhan, China, 2017*.
- Zhu, M., & Gupta, S. (2018, April). To prune, or not to prune: Exploring the efficacy of pruning for model compression. In *6th international conference on learning representations (ICLR), Vancouver, B.C., Canada, 2018*.
- Zoph, B., Vasudevan, V., Shlens, J., & Le, Q. (2018, June). Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition, Salt Lake City, UT, USA, 2018*.
- Zou, Q., Ni, L., Zhang, T., & Wang, Q. (2015). Deep learning based feature selection for remote sensing scene classification. *IEEE Geoscience and Remote Sensing Letters*, 12(11), 2321–2325. <http://dx.doi.org/10.1109/LGRS.2015.2475299>.