



HAL
open science

Neural Architecture Search in operational context: a remote sensing case-study

Anthony Cazasnoves, Pierre-Antoine Ganaye, Kévin Sanchis, Tugdual Ceillier

► To cite this version:

Anthony Cazasnoves, Pierre-Antoine Ganaye, Kévin Sanchis, Tugdual Ceillier. Neural Architecture Search in operational context: a remote sensing case-study. [Research Report] Preligens. 2021. hal-03343944

HAL Id: hal-03343944

<https://hal.science/hal-03343944>

Submitted on 14 Sep 2021

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

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

Neural Architecture Search in operational context: a remote sensing case-study

Anthony Cazasnoves ^{*}, Pierre-Antoine Ganaye ^{*}, Kévin Sanchis ^{*}, and Tugdual Ceillier

Preligens (ex-Earthcube), Paris, France
www.preligens.com, [name].[surname]@preligens.com

Abstract. Deep learning has become in recent years a cornerstone tool fueling key innovations in the industry, such as autonomous driving. To attain good performances, the neural network architecture used for a given application must be chosen with care. These architectures are often handcrafted and therefore prone to human biases and sub-optimal selection. Neural Architecture Search (NAS) is a framework introduced to mitigate such risks by jointly optimizing the network architectures and its weights. Albeit its novelty, it was applied on complex tasks with significant results – e.g. semantic image segmentation. In this technical paper, we aim to evaluate its ability to tackle a challenging operational task: semantic segmentation of objects of interest in satellite imagery. Designing a NAS framework is not trivial and has strong dependencies to hardware constraints. We therefore motivate our NAS approach selection and provide corresponding implementation details. We also present novel ideas to carry out other such use-case studies.

Keywords: Neural Architecture Search (NAS) · High Performance Computing (HPC) · Remote Sensing · Semantic segmentation.

1 Introduction

Deep learning has nowadays established itself as the reference framework to tackle complex signal processing problems. Their high level of performance and strong versatility make neural networks (NN) both best contenders on many public benchmarks and key enablers for multiple real-world applications, among which autonomous driving, expert translation system, 3D pointcloud processing or remote sensing. Reaching operational performances through deep learning is nonetheless not trivial. The quantity of high-quality labelled data is obviously one of the process cornerstones. Another one lies in the architecture of the network that will be trained.

In terms of architecture design, one can think of two potentially antagonist objectives to reach: finding the best performing structure possible while enforcing a complexity level suitable for training. This trade-off performance / simplicity is perceptible in the literature. When considering a specific task, some deceptively simple structures manage to reach very good performance levels and further gains can only be made at considerable complexity increase costs. Compare, for example, the UNet[1] to the DeepLab V3[2] when dealing with semantic segmentation. In most contributions, the network structure is purely handcrafted, either inspired by other *NN* or distilled know-how from non machine-learning algorithms. Many atypical architectures are, in this way, never probed: human design is subject to multiple bias – e.g. preference for symmetries in structures – that can lead to the exclusion of simple architecture exhibiting good performances.

^{*} Authors contributed equally to this paper

Automated Machine-Learning (Auto-ML)[3] is the most encompassing framework to tackle such concerns; the idea there being to model both the algorithm task and all your possible choices for the network architecture as optimization problems. Under such paradigm, these critical selections are performed through a much more complex gradient descent. Neural Architecture Search (NAS) is a sub-domain of Auto-ML: it consists in jointly optimizing the NN weights and its architecture to find the most suitable solution to a given problem – all other hyper-parameters remaining hand picked. While mainly applied on *lower complexity* tasks – e.g. image classification with large inter-classes variance, such as CIFAR-10[4] or ImageNet[5] – NAS has recently shown encouraging results on tasks closer to real-world use-cases. AutoDeepLab [6] for example managed to outperform SOTA results on the CityScape benchmark[7], thus proposing an interesting architecture for autonomous-driving applications.

In this technical paper, we evaluate the potential of NAS for one of our company use-cases: the semantic segmentation of objects of interest in satellite imagery. To the best of our knowledge, NAS was never applied in this operating context. We have 3 objectives. First, we aim to evaluate if this technology can prove useful in finding relevant network architectures for remote sensing applications. Considering the higher level of complexity involved in implementing a NAS framework, we also strive to provide a detailed technical feedback relevant to people interested in developing one. Finally, we intend to highlight the limitations encountered and introduce some ideas for further studies.

This paper is organized as follows. Section 2 first introduces the NAS problem formulation and the two main families of approaches used to perform such research. Selection of relevant methods is then motivated in Section 3. Experiments taxonomy is presented in Section 4 while results are detailed in Section 5. Conclusions and discussions of this study can be found in Section 6. Implementation challenges and related caveats are available in Appendix A.

2 Problem statement and related works

In this section, we'll briefly introduce how the NAS problem is often framed in literature before presenting the two main families of approaches to conduct said search.

2.1 Problem statement: an overview

We here follow the insight developed in [8, 9] and therefore encourage interested readers to consult these surveys for more in-depth details.

As aforementioned, the bread and butter of NAS is to automatically probe for the best architecture to tackle a specific problem. In supervised learning, a problem is mainly modeled through an annotated dataset. As such, letting \mathcal{D} be the space of all annotated datasets, \mathcal{A} the search space of NN architectures and \mathcal{M} the trained models one, the NAS search application Γ is defined under

$$\Gamma : \mathcal{D} \times \mathcal{A} \rightarrow \mathcal{M}. \quad (1)$$

In deep learning, upon selecting a specific architecture $a \in \mathcal{A}$ and addressing the problem considered $d = (d_{train}, d_{valid}) \in \mathcal{D}^2$, one wants to find the best set of network weights w by minimizing

$$\Gamma(a, d_{train}) = \arg \min_{m_{a,w} \in \mathcal{M}} \mathcal{L}(m_{a,w}, d_{train}) + \mathcal{R}(w), \quad (2)$$

where $m_{a,w}$ represents the trained model, \mathcal{L} the selected loss function and \mathcal{R} the regularization term.

Building upon this, NAS strives to find the best architecture a^* to handle problem d i.e. the one searched on d_{train} exhibiting best performance on d_{valid} according to the objective function \mathcal{I} :

$$a^* = \arg \max_{a \in \mathcal{A}} \mathcal{I}(\Gamma(a, d_{train}), d_{valid}). \quad (3)$$

\mathcal{I} being a black-box function, no direct optimization can be performed. Two families of methods have been proposed to get around this problem, it will be the topic of the following sections.

Before that, some concepts common to both families must be introduced. Enabling an automated search of architecture requires to define and model the search space \mathcal{A} embedding all possible architectures. One can for example choose to encode it as two complementary structures – see Figure 1. Firstly, a directed acyclic graph (DAG) that can encode all available networks topologies by:

- mapping tensors to its nodes
- mapping operations to its edges
- letting the graph direction follow the computation ones

Secondly, the set of all available operations \mathcal{O} that can be selected to perform said computations. In doing so, the search strategy can then select any topologies available in the DAG, pick any operations available in \mathcal{O} and instantiate a NN that will be trained and evaluated.




Fig. 1: A NN can be symbolically encoded using two structures: a directed acyclic graph (DAG) for its topology (left) and a dictionary of operations available to perform computations (right).

2.2 Discrete Formalism

The discrete approaches were historically the first propositions in the NAS community. These methods proceed in an exploratory way, that could be summed-up in three main phases:

1. choose a subset of NN (by picking topologies and related operations)
2. train this subset (classical deep learning)
3. evaluate it according to relevant criteria

As a first naive approach, one can of course iterate over these three steps by picking at random the probed subsets and recording the history of resulting performances to later choose its best performer. However, considering the combinatorial number of possible architectures, the time required to find a good contender for a given real-world problem – without supervision in this search – will be considerable. As we do have access to a relevant evaluation metric, it can be used to supervise the search. Two branches of algorithms are well suited for such a task: evolutionary algorithms (EA) and reinforcement learning (RL).

From a very general perspective, RL is based on a feedback loop mechanism. An agent performs actions affecting its environment, which impacts are then evaluated by a controller that provides the adequate guiding signal for the update policy. Iterating in this way, the agent adapts and improves its performances – i.e. its ability to correctly behave on the training task. Following this idea in NAS, one can choose to set a controller acting as architecture sampler, guided by its achieved performance. In [10] a RNN acts as the controller that provides a variable-length token encoding the NN architecture. The NN performances are used as rewards driving the RNN – either through REINFORCE loss [10] or Proximal Policy Optimization [11]. Other contributions chose to exploit the Q-Learning paradigm to perform this sampling [12, 13].

EA strives to emulate a nature-inspired adaptation processed, as framed in Darwinian theory. Using a specific fitness criterion, and a population of potential solutions, the algorithm loops through the following steps to select the best candidates:

1. best specimens selection (information selection)
2. matting (information sharing)
3. mutating (random information perturbation)

There is a large variety of ways to perform these steps and to handle specific inherent risks like sub-optimal elitist selection. In NAS, the population will be composed of architectures genotypes – structures encoded as strings, for instance. The evaluation will be performed on their phenotypes, i.e. NN that have been instantiated and trained. In terms of mutations, two types can be explored jointly: operands switch and topological modifications. [14] uses tournament selection while [15] builds upon it by introducing a new parameter: ageing. In a running population, a fraction of the oldest individuals is removed at each cycle – indiscriminately from their performance – to limit the risk of early stagnation and provide a wider exploration breadth. An interesting proposal is introduced in [16]: maintaining an architecture population by sampling the Pareto front performances / resources requirement. It is shown that along generations, the achieved trade-offs are improving significantly.

All the techniques cited above suffer from the same major limitation: no matter how smart the exploration strategy, trying to discover a^* in this way amounts to an almost brute-force exploration of all available architectures. One can take advantage of an HPC infrastructure to improve the chances: at each iteration, the number of candidates will scale in agreement with the number of GPUs available. To have a real impact, however, the resources involved are colossal – as seen in many papers, they are on the scale of hundreds to thousands of GPUs. Recent studies [9, 17, 18] even point out that EA and RL may not outperform random search. These drawbacks have led to the development of another paradigm: the idea that a better problem formulation could partially alleviate the hardware constraint and provide a greater chance of success in finding a good architecture.

2.3 Continuous Formalism

Building on the idea of simplifying the training process and reducing the significant computation cost of aforementioned methods, a new set of approaches based on the training of

a single, dense network for architecture search were proposed. Herein, we review the main methods of this group and define the formalism which they are based upon.

DARTS [19] is the first method to consider architecture search as the training of a single differentiable hyper-network that contains the whole operation search space, using the principle of continuous relaxation. We review this method as the core building block of continuous approaches in architecture search. As mentioned previously, \mathcal{O} is our operation search space, a set of predefined operations that can be applied to a tensor x , such as a convolution or a pooling. DARTS aims at building a search network \mathcal{S} , that is a directed acyclic graph (DAG) formed of a single search cell that is used several times. This search cell is in fact another DAG: training \mathcal{S} amounts to assigning an operation o_i from \mathcal{O} to each edge of this sub graph.

Formally, during training, each edge of the search cell is assigned to a mixed operation o_m defined as a weighted sum over all operations from \mathcal{O} :

$$o_m^{(i,j)}(x) = \sum_{o \in \mathcal{O}} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in \mathcal{O}} \exp(\alpha_{o'}^{(i,j)})} o(x) \quad (4)$$

Where $\alpha_o^{(i,j)}$ is a vector of dimension $|\mathcal{O}|$ associated with an edge (i, j) , that is jointly learned alongside operation weights w with gradient descent.

Note that, as an edge (i, j) connects two nodes n_i and n_j , a node’s output can be expressed as follows:

$$n_j = \sum_{i < j} o_m^{(i,j)} \quad (5)$$

Once training of w and α_o is done, a discrete architecture can be decoded from \mathcal{S} by replacing each mixed operation o_m with the best operation o^* w.r.t α , defined as:

$$o^{*(i,j)} = \operatorname{argmax}_{o \in \mathcal{O}} \alpha_o^{(i,j)} \quad (6)$$

Building on the concept of differentiable architecture search, several methods were proposed afterwards, which resulted in a new family of approaches based on continuous relaxation, as opposed to the more traditional discrete ones.

[20] proposed an incremental improvement to solve a well identified problem in DARTS-based approaches, where sometimes skip or any other operation that contains few or no parameters would be highly preferred over operations with parameters (e.g convolutions), leading to poor performance once the decoded discrete network is retrained. This work suggests to use manually defined criteria to control the number of skip connections in the discretized search cell. [21] tackled the issue of memory efficiency induced by DARTS-based approaches, inherent to the construction of a super-network that encapsulates the whole operation search space (e.g mixed operations). By using channel-wise partially connected operations, this method manages to significantly reduce the memory footprint of the search network. Additionally, PC-DARTS introduces edge normalization, a new form of regularization that helps overcoming the *skip-connection problem* described above.[6] goes a step

further by learning and decoding both the network structure and the cell structure simultaneously. It proceeds by building a densely connected search network – as opposed to cell-based networks traditionally used by DARTS approaches – that is also trained using continuous relaxation and additional β parameters encoding the network connection probabilities between cells. Finally, the Viterbi algorithm is used to decode a network structure, as in this case decoding an optimal structure from a densely connected DAG means finding an optimal path from the network’s start to its end.

3 Selected methods

Considering the dedicated literature, multiple methods can be used to tackle our use-case through NAS. We will here introduce the reasons behind our methods and implementation choices.

3.1 Choosing the most appropriate paradigm

As mentioned before, Neural Architecture Search algorithms can be separated into two categories: discrete search and continuous relaxation of the search space. Although we explored both paradigms, we determined that one was more adapted to our use case. While exploring the typical discrete framework by training multiple architectures using a simple genetic algorithm, we encountered several challenges, the most prominent one being the unpredictable memory footprint and training time of such approaches. More specifically, we could not anticipate the GPU RAM usage of the sampled architectures, leading to uncontrolled program failures due to memory allocation limit[22].

Considering this observation, we explored the more stable continuous approaches and deemed them adapted to our use-case. We first considered the classical DARTS approach. While this method offers a significant reduction in memory and training time, it is constructed on a cell-based network and tailored to the image classification problem. In order to build suitable architectures for our image segmentation task, we replaced the usual cell-based structure by another one, similar to U-Net [1]. Undoubtedly, this new structure is not as lightweight as the original one, and thus lead to an increase in memory consumption. However, this consumption is constant throughout training and methods like the ones used for PC-DARTS [21] can help reducing it.

3.2 Reducing the memory footprint

In the purpose of dealing with the memory issue described in the previous section, we explored [21] as a potential solution. *Partial channel connection* [21] is an effective technique to reduce the memory overhead of computing mixed operations inside every search cell. For a given edge (i, j) containing a mixed operation $o_m^{(i,j)}$, the input tensor x is split into two parts along the channel axis. The computation is only applied on the first part, while the second part is simply propagated and left unchanged using a skip connection. Formally, a mask $M_{i,j}$ assigns 1 to channels inside the first part, and 0 to the others. Thus, $o_m^{(i,j)}$ is defined as:

$$o_m^{(i,j)}(x; M_{i,j}) = \sum_{o \in \mathcal{O}} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in \mathcal{O}} \exp(\alpha_{o'}^{(i,j)})} o(M_{i,j} * x_i) + (1 - M_{i,j}) * x_i \quad (7)$$

The weighted sum is only applied on a subset of channels – as opposed to Equation 4 where every operation is applied on x – which reduces the overall memory footprint. Practically, the mask is applied by randomly sampling a portion of $1 / K$ channels that will be processed – K being an hyper-parameter. This solution creates an adjustable trade-off between efficiency, when K is large, and accuracy, when K is small. Additionally, sampling a subset of channels adds a form of regularization to the architecture search, as most channels are skipped and the influence of α on the overall operation selection is reduced. However, random sampling also adds instability to the training process.

We follow [21] and use *edge normalization* in conjunction with partial channel connection. Formally, edge normalization is defined as the addition of a value, denoted $\gamma^{(i,k)}$, that is bound to each edge of a search cell, so that the expression of a node n_j (5) becomes:

$$n_j = \sum_{i < j} \frac{\exp(\gamma^{(i,j)})}{\sum_{i' < j} \exp(\gamma^{(i',j)})} o_m^{(i,j)} \quad (8)$$

and the decoding rule (6) is now expressed as:

$$o^{*(i,j)} = \operatorname{argmax}_{o \in \mathcal{O}} \alpha_o^{(i,j)} \gamma^{(i,j)} \quad (9)$$

Similar to α , γ are learned throughout the training and bring stability to counteract the random, unstable nature of channel sampling. Albeit the use of partial channel connection and edge normalization solved the memory footprint overhead inherent to DARTS-based approaches, we empirically observed an additional issue bound to the operation selection step, detailed in the following section.

3.3 Improving architecture selection

While the continuous relaxation scheme has gained a lot of interest for its efficiency, we experimentally observed several failure cases: poor operations diversity within the decoded cell (see Figure 2) and performance discrepancy between the search step and the retraining steps. These problems have already been mentioned in section 3 and in [23].




Fig. 2: Graph of a decoded (post-search) DARTS cell. Only operations without learnable parameters were selected. This type of weight-free architecture is called degenerated because of its lack of modeling capacity.

To correct the generalization issues, [24] formulated an optimization method called GAEA (Geometry Aware Exponentiated Algorithm), which is more suitable than Stochastic Gradient Descent (SGD) for updating the architecture’s parameters. Indeed, GAEA is based on exponentiated gradient descent which is known to make convergence faster for constrained parameters and ultimately provides sparsity properties that are helpful to reduce the generalization gap induced by DARTS [19]. In the optimization method proposed in [24], the network’s parameters are still learned by SGD, while GAEA is used to optimize only θ , the architecture’s weights.

GAEA can be plugged into continuous relaxation-based approaches such as PC-DARTS by substituting the learned architecture parameters $\alpha_o^{(i,j)}$ by $\theta_o^{(i,j)} \in \Theta = \mathbb{R}^{|\mathcal{O}|}$. The main difference between $\alpha_o^{(i,j)}$ and $\theta_o^{(i,j)}$ is that $\theta_o^{(i,j)}$ is constrained to be a probability vector (probability of operation $o \in \mathcal{O}$ for edge (i, j)), where $\alpha_o^{(i,j)}$ is only considered as a logit that is used by the softmax layer. Optimizing $\theta_o^{(i,j)}$ consists in two steps: updating with exponentiated gradient (equation 10), then re-parameterizing $\theta_o^{(i,j)}$ (equation 11):

$$\theta^{(i,j)} \leftarrow \theta^{(i,j)} \odot \exp \left(-\eta \nabla_{\theta} f \left(\mathbf{w}, \theta^{(i,j)} \right) \right) \quad (10)$$

$$\theta_o^{(i,j)} \leftarrow \frac{\theta_o^{(i,j)}}{\sum_{o' \in \mathcal{O}} \theta_{o'}^{(i,j)}} \forall o \in \mathcal{O} \quad (11)$$

where η is the learning rate and $\theta_o^{(i,j)}$ is the probability of operation o being selected for edge (i, j) at the next parameter update.

One way to measure the progression of architecture search is to use the entropy of θ (the lower, the more certain the choice of operation o is). The active learning community uses entropy as a measure of uncertainty. Here, it can be applied on θ to derive a convergence indicator for the architecture’s parameters. As mentioned previously, the main advantage of GAEA is its improved convergence speed, which translates into an increased θ sparsity and thus lower entropy. When compared to PC-DARTS, we can observe in Figure 3 that GAEA reaches a lower entropy by a significant margin.




Fig. 3: Evolution of the entropy of θ during architecture search for PC-DARTS [21] and GAEA [24]. Figure from [24].

4 Experiments

4.1 Use-case presentation

We perform neural architecture search for a semantic segmentation task. Our dataset contains around 23,000 patches with 12,000 observables, tiled from *Maxar/DigitalGlobe* satellites WorldView 1, 2, and 3 and GeoEye 1 publicly available full-size images. The patches have a dimension of 512×512 pixels, with an overlap of 180 pixels between adjacent patches. We use 18,000 patches for training and 5,000 patches for validation. We do not follow the common 50/50 split of the training dataset used by DARTS [19] for optimizing layer weights and architecture weights, but instead follow the method used by GAEA [24], which consists in duplicating and using the full training dataset for both the layer weights and the architecture weights. Finally, we evaluate our decoded architectures on a testing set composed of 38 full-size images, containing 906 observables. The objective of our task is to assign one of two classes to each pixel: *background* or *helicopter*.




Fig. 4: Typical examples from our dataset. Helicopters can be found at any position and have different forms. Some images have a significant class imbalance or contain no objects. Sometimes, helicopters can be hard to distinguish from the background.

Remote sensing images bring various challenges. First, there is usually a significant class imbalance between background and objects, as objects are relatively small. Objects also don't follow any placing rules and can be found at any position on a patch. Finally, tiling an image also creates corner cases where objects are at the edge of a patch, sometimes partially visible, because of cuts between two patches. Additionally, in our specific case, objects can have different shapes, as illustrated in Figure 4. To solve this task, we explore different search spaces, network and cell structures thoroughly described in the following section.

4.2 Explored topologies

We present all conducted experiments sorted by *topology*. We define a topology \mathcal{T} as a unique combination of network structure \mathcal{S}_n , cell structure \mathcal{S}_c , and operation search space \mathcal{O} . Presented topologies adopt the following naming convention: $\mathcal{S}_c\text{-}\mathcal{S}_n$. To perform these experiments, we developed an internal NAS framework that aims to be as generic as possible. Further details about this framework are available in **Appendix A**.

We first review the **ResNeXt-UNet** topology. In this setting, we use a network structure identical to U-Net [1], more fitted for instance segmentation than the classical cell-based structure of DARTS. Furthermore, the DARTS cell is discarded in favor of a structure similar to a ResNeXt [25] block. The reasoning behind the choice of this topology is to restrain our search network \mathcal{S} to follow already proven structures on the task we tackle. In other words, we apply NAS on a well-defined setting and try to incrementally improve upon an already convenient structure, as well as evaluating the method on a simpler case (i.e with less degrees of freedom) than DARTS.

In [19, 21, 20], the usual operation search space that we denote \mathcal{O}_{base} contains the following operations: 3×3 and 5×5 dilated convolutions, 3×3 and 5×5 depthwise separable convolutions, 3×3 average pooling, 3×3 max pooling, skip connection, and zero (e.g multiplying the input tensor by 0). To compensate for the lower degrees of freedom, we design a larger operation search space \mathcal{O}_{large} containing: 3×3 , 5×5 and 7×7 convolutions, 3×3 , 5×5 and 7×7 depthwise convolutions, 3×3 , 5×5 and 7×7 spatial separable convolutions, skip connection, and zero. In our implementation, we respectively call the operations *conv2d_1*, *conv2d_2*, *conv2d_3*, *depthconv2d_1*, *depthconv2d_2*, *depthconv2d_3*, *splitconv2d_1*, *splitconv2d_2*, *splitconv2d_3*, *skip*, and *cut*. Essentially, we replaced pooling operations with normal convolutions, and added a 7×7 kernel size for all convolutions. The ResNeXt-UNet topology is depicted in Figure 5.




Fig. 5: The ResNeXt-UNet topology. The upper graph is the search network, the lower graph is the search cell.

Concurrently, we explore the **DARTS-UNet** topology, the closest to a typical DARTS-based topology. Building on the previous intuition, we use the same U-Net structure that we associate with a classical DARTS cell, depicted in Figure 6b. Similarly to [19, 21, 20], our search cell contains 4 blocks, and we use \mathcal{O}_{base} as an operation search space. We also follow the same connection scheme as a cell-based network structure, where each search cell is connected to its parent node (i.e the closest first order neighbour in topological order), and its grand-parent node (i.e the closest second order neighbour in topological order).

While decoding the trained DARTS cell, we keep the \mathcal{K} strongest input edges per block based on:

$$\alpha_{o^*}^{(i,j)} = \max_{o \in \mathcal{O}} \alpha_o^{(i,j)} \quad (12)$$

and discard the others. In other words, we sort edges using the α value associated with their selected operation and retain the \mathcal{K} largest values for each block. Akin to [19, 21, 20], we set \mathcal{K} to 2, so each block is connected to at most 2 input edges after decoding.

Using this topology, we aim at providing an adaptation of DARTS-based approaches for an instance segmentation task, only replacing the network structure. To the best of our knowledge, there is no existing work that propose such a direct application of DARTS on image segmentation. Albeit some methods such as [26, 6] tackle this task, they require significant changes and are not straight applications of DARTS. Building on the idea of a two-step decoding from [6], we additionally consider a denser network structure in the **DARTS-UNet++** topology. Based on [27], this structure remains close to U-Net, but is also dense enough to contain novel encoder-decoder structures.

In [6], the cell-based network from DARTS is replaced with a densely connected one that is trained similarly to search cells, i.e with β continuous values associated with all edges. After training, a single path is decoded using the Viterbi Algorithm, where softmaxed β values are considered as transition probabilities. This optimal path is then used as an encoder and attached to a fixed decoder. We extend this idea to train a dense encoder-decoder search network. We define a network structure identical to [27], as shown in Figure 6a, and aim at discovering an encoder-decoder structure. We associate this network structure with DARTS search cells. As we are looking for an encoder-decoder network, decoding a single path is not suitable. Instead, we propose a simple, greedy multi-path algorithm that computes the average log-likelihood of all possible paths. While this algorithm is not optimized, the set of paths is limited enough – e.g. lesser than 100 – to remain computationally tractable. Upon gathering their scores, one still needs to select a relevant subset of \mathcal{K} paths used in the decoded architecture. \mathcal{K} may either be considered as a fixed hyper-parameter, or one can try to build a statistical rule to decide the adequate cardinal to retain. We attempted a simple tryout that exhibited convincing properties in our experiments. Letting $p_{sc} \in \mathbb{R}^{|p_{sc}|}$ be the set of paths scores, one can compute its corresponding mean μ and variance σ . Assuming Gaussian nature, \mathcal{K} may therefore be determined as:

$$\mathcal{K} = |p_{sc}, p_{sc} \geq \mu + 3 * \sigma| \quad (13)$$

We describe the results of our experiments in the next section.

5 Results and observations

In this section, we present and discuss the results of our experiments. We compare the performance of decoded and retrained search networks with a custom *baseline network* following a ResNet-UNet architecture – a U-Net where encoder blocks are replaced with ResNet [28] blocks, and tailored for our dataset. In other words, hyperparameters such as the number of filters of each convolutional layer, or the number of layers in each block, are handpicked to achieve the best possible results on our dataset. We report in the following sections the *Precision-Recall curves* on the testing set and evaluate our topologies against the baseline network.

For all experiments described in the following subsections, similar to [19], we use an initial number of channels of 16 and start training α values at epoch 15, using GAEA [24] as an optimizer for α , and SGD with momentum for layer weights w . Furthermore, we use a cosine annealing learning rate with an initial value of 0.01, down to 0 for SGD, and a fixed learning rate of 0.1 for GAEA. Finally, we apply a weight decay of 10^{-3} . Furthermore, similar to the

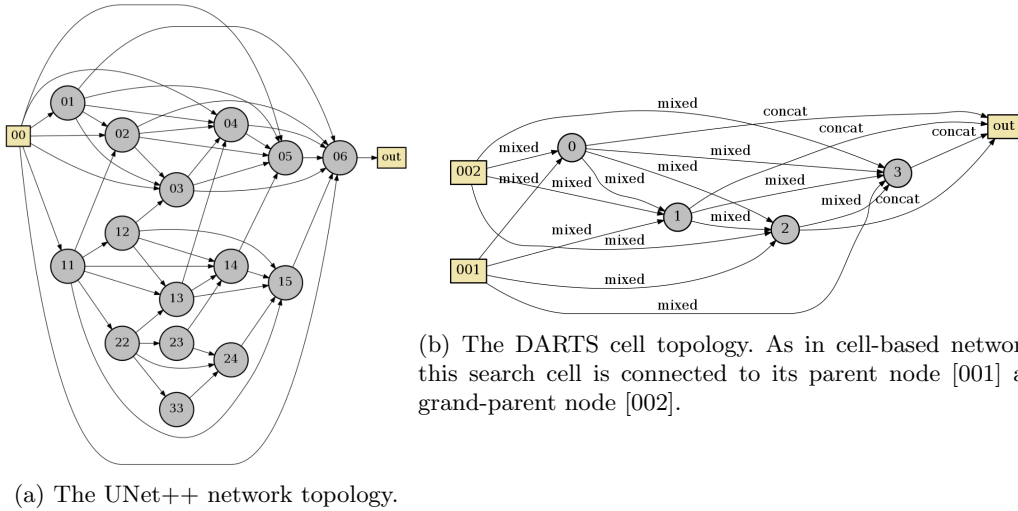


Fig. 6: Topologies of the UNet++ network (a) and the DARTS cell (b)

search step, we use SGD with the cosine annealing rate mentioned above as an optimizer for the training of the decoded architecture. Finally, we use weighted cross entropy as a loss for all experiments.

5.1 ResNeXt-UNet

We first investigate the results of the ResNeXt-UNet topology. We train the search network for 50 epochs, using a batch size of 2, a partial channel connection of 1/8, and no edge normalization as it is not applicable for a ResNeXt cell, since each block is only connected to at most one input edge. Once a network architecture is decoded, we train it for 80 epochs and select the checkpoint that achieves the best *MeanIoU* on the validation set, starting from the 30th epoch.

We report the achieved *Precision-Recall curve* in Figure 7. Although we did not manage to do better than the baseline network, we emphasize the relatively small difference between both approaches and compare time and efforts required to elaborate them. As the baseline network is manually designed and optimized for our dataset, it requires a significant amount of time and efforts to create, whereas our ResNeXt-UNet only required a single architecture search. Additionally, we did not increase the number of filters or add any form of regularization in the decoded ResNeXt-UNet, albeit the wide usage of such a trick in DARTS-based approaches [19, 21, 20, 6]. Finally, we did not use any form of data augmentation while retraining our decoded network, which is not true for the baseline.

In Figure 8, we show the decoded ResNeXt cell. We notice the majority of selected convolutions uses a 5×5 kernel, except *depthconv2d_3* which uses a 7×7 kernel. Furthermore, selected operations contain a good variety of convolutions, as all types are chosen twice, and only one skip connection, indicating that the training successfully avoided the *skip connection problem*. Finally, each tower branch is different – which may seem counter-intuitive – and contains no real logic. It is safe to assume it can be difficult for a human to manually design this architecture.

We present a small sample of predictions on the testing set in Figure 9. Globally, our network makes few false positives, but has a hard time contouring objects with precision, although being able to detect the general shape.




Fig. 7: Precision-Recall values as a function of IoU thresholds.




Fig. 8: The decoded ResNeXt cell (K=2).

5.2 DARTS-UNet

In DARTS [19], the authors experiment on image classification with two types of cells, namely the normal and reduction ones. Both share a common search space, the main difference being that the stride of the first nodes within the reduction cell is set to 2 (instead of 1), in order to apply a spatial reduction of the resolution. In the DARTS-UNet experiment, we follow the same scheme by designing a UNet [1] architecture with reduction cells within the encoder and normal cells with the decoder part.

The DARTS-UNet search space follows the one proposed in [21], it is only composed of pooling layers, separable and dilated convolutions and skip connection. Because these operations are light in terms of parameters, we can increase the batch size from 2 to 4, compared to other experiments. The training of the best architecture is done with SGD for 200 epochs, as we observed a slower convergence compared to the ResNeXt-UNet architecture.

During our experiments, the maximum IoU for the search step was 0.52, while the training of the sampled architecture led to an IoU of 0.73. The cell structures obtained after the search step can be visualised in figure 10. We notice the heavy presence of average poolings in both normal and decoded cells, and the absence of skip connections and max poolings.






Fig. 9: Predictions on the testing set. **(left)** The ground truth. **(right)** The prediction.

Furthermore, a good variety of convolutions in terms of type and kernel size can also be noticed. Once again, the absence of skip connections may indicate that the training did not fall into the *skip connection problem*, despite preferring operations with a low amount of parameters (e.g. pooling). The fair amount of convolutions seems to indicate a successful training nonetheless. Finally, we notice a lower amount of average pooling inside the reduction cell, which seems intuitive since this type of cell already reduces the number of parameters by dividing the input feature map by two, and thus does not need much pooling on top of it.



(a) DARTS-UNet normal cell decoded after search.



(b) DARTS-UNet reduction cell decoded after search

Fig. 10: DARTS-UNet cells obtained after the architecture search.

5.3 DARTS-UNet++

For this topology, the search step is performed on 50 epochs, and the decoded architecture is trained on 80 epochs. As main differences, edge normalization is here applied and we

use two stem-cells – fixed cells used to reduce the input feature map’s size at the start of the network – before entering the main DARTS-UNet++ architecture – thus reducing the spatial resolution to 128×128 pixels. This stemming scheme is mainly due to VRAM consumption limitation and enables us to perform a wider search in the main architecture section, as the DARTS-UNet++ is significantly heavier than other architectures in terms of memory requirements.

As presented in Figure 11, the decoded network exhibits very dense connectivity at maximum spatial resolution and retains the usual U-Net high semantic path. This intuitively makes sense: to perform proper spatial segmentation, processing full resolution patches is critical while the classification task can benefit from higher semantic content. The decoded cell displays a variety of operations with both convolution layers and pooling/skip layers. For this experiment, only classical convolutions were ultimately chosen however.




Fig. 11: Decoded UNet++ network (left, colored paths) and corresponding DARTS cell (right). Compared to a standard UNet, the network displays a dense connectivity at higher spatial resolution.

On a qualitative basis – see Figure 12 – the performed segmentations delineate quite well the shapes of the detected helicopters. However, from a quantitative standpoint, this architecture didn’t manage to outperform either our baseline or the structure found using the ResNeXt-UNet topology.

5.4 Random search baselines

Evaluating the potential of a new method always requires to compare against a competitive alternative. Following recommendations made on NAS experiment reproducibility [29], we decided to compare our best searched DARTS cells against random search. To this end, based on the DARTS-UNet and ResNeXt-UNet cell topologies, we apply a random selection of each node’s operation by uniformly sampling with replacement. We call these random cells as baselines in our experimental protocol.




Fig. 12: From left to right: image, ground truth and DARTS-UNet++ predictions.

Random search is a simple alternative to DARTS, though it is still necessary to train the sampled cell in the same way as before to obtain a fair comparison. In figure 13 we illustrate the maximum IoU obtained after the training of the selected cells (through DARTS or random search). We experimentally observed that random search can reach performances similar to DARTS, thus using DARTS-based NAS on this specific use case amounts to finding a best performing cell with random search.




Fig. 13: Performance comparison of random search and DARTS for two types of cell topology.

6 Conclusions and feedback

As mentioned in this paper opening, the aim of this first study was threefold. NAS being presented as an interesting network design tool in the literature, our first intent was to find the most relevant methods to build such a tool and to implement a framework based on it. Upon building this tool, our second objective was to put it to use on one of our company use-case and benchmark whether or not more suitable architectures could be found using it. Finally, and following these experiments, we wanted to present some first conclusions and use our feedback to propose new ideas to build on this study.

In terms of theory and method, NAS is still a very young subject, thus lacking stability and in depth insights. When compared to classical deep learning, the number of hyper-parameters drastically increases, and when considering the continuous relaxation framework, the joint optimization to perform is vastly more complex. A striking example of how this complexity translates to experiment instability: the high variability in terms of network structure and the corresponding performance between two similar searches. Indeed the probed architectures are initialized at random – as we obviously can’t rely on usual ImageNet weights – and therefore suffers from their higher number of trained parameters. This lack of maturity is also exhibited by the numerous studies arguing that random search may have similar performances – or better ones – given similar search time. This is one of the main conclusions of this study, which wasn’t especially expected when considering performances displayed by the technology in some of the NAS key papers.

As noted in the literature, the computing power required to use such techniques is still a very limiting factor – even when considering recent techniques to reduce it. Not every company can afford ownership or access to an HPC infrastructure. Moreover, implementing a framework running properly in such multi-nodes/multi-GPUs context is not trivial. Details are presented in Appendix A, but in a nutshell, one ends up using very experimental parts of relevant backend and libraries and therefore encounters rare cases of failures and unexpected behaviours – for which only a very limited community may be able to help.

Applying NAS to one of our use-cases, we obtained mixed results. On the one hand, when challenging an heavily tailored architecture, we did not manage to find an outperforming structure. Multiple reasons exist for this outcome – e.g. we searched for a single cell structure when the reference architecture possess various tailored ones. On the other hand, when considering the sheer time involved in designing our production architecture, NAS managed to find a competitive structure much faster and at a far lower human-involvement cost – upon completion of the framework implementation. Taking into account the possible biases in our helicopter experiment, we benchmarked NAS against a less use-case tailored architecture. Preliminary results are more encouraging: the NAS-found structures manage to reach the same level of performance – and even to outperform it. However, we would not use NAS to design a better NN for this use-case.

By conducting this first study, we managed to get a better grasp on the core theory and technology involved in NAS and implement a first framework enabling us to conduct our experiments. Following these results, further studies will focus on potentially missing tools to our framework or better conditioning of our search strategy. As already noted, a large randomness is introduced in our probed super-network structures due to the increased number of parameters and inability to rely on pre-trained weights. As such, a tool to assess structural soundness, previous to any search / training phase, would be a great asset in NAS. Taking inspiration on recent propositions [30, 31, 32], one could investigate criteria such as *expressivity* and *trainability* to build such a metric. Regarding the high complexity involved when searching for an entire network / cell structure, one could try to reduce it and

embed it in an already well designed NN. As an example, in an handcrafted network, one could try to improve the global performance by NAS-searching for a very limited subsection of the network – e.g. at custom layer level. A question then arises: how to select the blocks that need improvement ? During search steps, we noticed that some large variation can exist between α entropy levels of the different cells composing the network. As such, the corresponding gradient applied during back-propagation could be more or less important, and our intuition is that such a criteria could be used to estimate which part of the network will benefit more from a change in its design.

Following this first study, it is indeed our intuition that one needs to better condition the propositions made in the literature to enable a really suitable use of NAS in industry applications.

Acknowledgements This work was granted access to the HPC resources of IDRIS under the allocation 2020-GC-101472 made by GENCI.

References

- [1] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. “U-Net: Convolutional Networks for Biomedical Image Segmentation”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (May 2015).
- [2] Liang-Chieh Chen et al. *Rethinking Atrous Convolution for Semantic Image Segmentation*. 2017. arXiv: 1706.05587 [cs.CV].
- [3] Quanming Yao et al. *Taking Human out of Learning Applications: A Survey on Automated Machine Learning*. 2019. arXiv: 1810.13306 [cs.AI].
- [4] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. “CIFAR-10 (Canadian Institute for Advanced Research)”. In: (). URL: <http://www.cs.toronto.edu/~kriz/cifar.html>.
- [5] Jia Deng et al. “Imagenet: A large-scale hierarchical image database”. In: *2009 IEEE conference on computer vision and pattern recognition*. Ieee. 2009, pp. 248–255.
- [6] Chenxi Liu et al. “Auto-DeepLab: Hierarchical Neural Architecture Search for Semantic Image Segmentation”. In: *CoRR* abs/1901.02985 (2019). arXiv: 1901.02985. URL: <http://arxiv.org/abs/1901.02985>.
- [7] Marius Cordts et al. “The Cityscapes Dataset for Semantic Urban Scene Understanding”. In: *CoRR* abs/1604.01685 (2016). arXiv: 1604.01685. URL: <http://arxiv.org/abs/1604.01685>.
- [8] Thomas Elsken, Jan Hendrik Metzen, Frank Hutter, et al. “Neural architecture search: A survey.” In: *J. Mach. Learn. Res.* 20:55 (2019), pp. 1–21.
- [9] Martin Wistuba, Amrith Rawat, and Tejaswini Pedapati. “A survey on neural architecture search”. In: *arXiv preprint arXiv:1905.01392* (2019).
- [10] Barret Zoph and Quoc V Le. “Neural architecture search with reinforcement learning”. In: *arXiv preprint arXiv:1611.01578* (2016).
- [11] Barret Zoph et al. “Learning transferable architectures for scalable image recognition”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, pp. 8697–8710.
- [12] Bowen Baker et al. “Designing neural network architectures using reinforcement learning”. In: *arXiv preprint arXiv:1611.02167* (2016).
- [13] Zhao Zhong et al. *Practical Block-wise Neural Network Architecture Generation*. 2018. arXiv: 1708.05552 [cs.CV].
- [14] Hanxiao Liu et al. *Hierarchical Representations for Efficient Architecture Search*. 2018. arXiv: 1711.00436 [cs.LG].
- [15] Esteban Real et al. *Regularized Evolution for Image Classifier Architecture Search*. 2019. arXiv: 1802.01548 [cs.NE].

- [16] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. *Efficient Multi-objective Neural Architecture Search via Lamarckian Evolution*. 2019. arXiv: 1804.09081 [stat.ML].
- [17] Yuqiao Liu et al. *A Survey on Evolutionary Neural Architecture Search*. 2021. arXiv: 2008.10937 [cs.NE].
- [18] Kaicheng Yu et al. *Evaluating the Search Phase of Neural Architecture Search*. 2019. arXiv: 1902.08142 [cs.LG].
- [19] Hanxiao Liu, Karen Simonyan, and Yiming Yang. “DARTS: Differentiable Architecture Search”. In: *International Conference on Learning Representations*. 2019. URL: <https://openreview.net/forum?id=S1eYHoC5FX>.
- [20] Hanwen Liang et al. *DARTS+: Improved Differentiable Architecture Search with Early Stopping*. 2020. arXiv: 1909.06035 [cs.CV].
- [21] Yuhui Xu et al. “PC-DARTS: Partial Channel Connections for Memory-Efficient Architecture Search”. In: *International Conference on Learning Representations*. 2020. URL: <https://openreview.net/forum?id=BJ1S634tPr>.
- [22] Yanjie Gao et al. “Estimating GPU Memory Consumption of Deep Learning Models”. In: *ESEC/FSE ’20*. The 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Industry Track. ACM, Nov. 2020, pp. 1342–1352. URL: <https://www.microsoft.com/en-us/research/publication/estimating-gpu-memory-consumption-of-deep-learning-models/>.
- [23] Arber Zela et al. “Understanding and Robustifying Differentiable Architecture Search”. In: *International Conference on Learning Representations*. 2020. URL: <https://openreview.net/forum?id=H1gDNyrKDS>.
- [24] Liam Li et al. “Geometry-Aware Gradient Algorithms for Neural Architecture Search”. In: *International Conference on Learning Representations*. 2021. URL: <https://openreview.net/forum?id=MUSYkd1hxRP>.
- [25] Saining Xie et al. “Aggregated Residual Transformations for Deep Neural Networks”. In: *CoRR* abs/1611.05431 (2016). arXiv: 1611.05431. URL: <http://arxiv.org/abs/1611.05431>.
- [26] Yu Weng et al. “NAS-Unet: Neural Architecture Search for Medical Image Segmentation”. In: *IEEE Access* 7 (2019), pp. 44247–44257. DOI: 10.1109/ACCESS.2019.2908991.
- [27] Zongwei Zhou et al. “Unet++: A nested u-net architecture for medical image segmentation”. In: *Deep learning in medical image analysis and multimodal learning for clinical decision support*. Springer, 2018, pp. 3–11.
- [28] Kaiming He et al. “Deep Residual Learning for Image Recognition”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2016.
- [29] Liam Li and Ameet Talwalkar. *Random Search and Reproducibility for Neural Architecture Search*. 2019. arXiv: 1902.07638 [cs.LG].
- [30] Joseph Mellor et al. *Neural Architecture Search without Training*. 2021. arXiv: 2006.04647 [cs.LG].
- [31] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. *Neural Architecture Search on ImageNet in Four GPU Hours: A Theoretically Inspired Perspective*. 2021. arXiv: 2102.11535 [cs.CV].
- [32] Ming Lin et al. *Zen-NAS: A Zero-Shot NAS for High-Performance Deep Image Recognition*. 2021. arXiv: 2102.01063 [cs.CV].
- [33] Message P Forum. *MPI: A Message-Passing Interface Standard*. Tech. rep. USA, 1994.
- [34] Lisandro D Dalcin et al. “Parallel distributed computing using Python”. In: *Advances in Water Resources* 34.9 (2011), pp. 1124–1139.
- [35] Martíen Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: <https://www.tensorflow.org/>.

- [36] Alexander Sergeev and Mike Del Balso. “Horovod: fast and easy distributed deep learning in TensorFlow”. In: *arXiv preprint arXiv:1802.05799* (2018).

A Implementation for HPC use

Thanks to the *Grand Challenge Jean-Zay* program, initiated by the GENCI, we were allocated computing hours on their HPC infrastructure. This enabled us to conduct NAS experiments that would otherwise not have been possible on our proprietary hardware. In terms of computing power, the partition of the cluster opened to us owned 351 nodes each equipped with 4 GPUs NVIDIA V100, 2 Intel cascade processors and 192 Gb of RAM.

To fully take advantage of these hardware capabilities, our framework needed to properly handle both model parallelism (discrete approach) and data parallelism (continuous approach). In the following, we will develop on the tools we used and the issues we encountered during our implementation.

A.1 Model parallelism: Mpi4Py

Although not mentioned in this paper, we conducted experiments using evolutionary NAS (E-NAS). To do so, we chose to implement the search using a multiprocessing dispatch approach:

- a parent process acting as an orchestrator
- children processes performing training / evaluation

By *orchestrator*, we here mean that the process handles the architectures population evolution – of course, to prevent idling, it also performs training tasks. The Message Passing Interface (MPI)[33] is a library widely used on HPC infrastructures. It was the perfect dispatching tool for our orchestrator and, as our framework was coded in Python, we relied on the dedicated package MPI4Py [34].

For clarity purpose, the MPI dispatching pseudo-code is presented in Algorithm 1. In this implementation, upon architectures dispatch to all GPUs by the orchestrator, we ensure that each device is indeed used and training a specific NN. By nature, this code section is asynchronous: although computations are performed on an homogeneous infrastructure, each NN architecture will need a different time to be trained on a specific number of epochs – relying heavily on how complex the probed NN architecture is. Synchronisation therefore is ensured by the orchestrator: computations are dispatched by it and, upon completing its own training, it awaits a signal from each GPU providing the results of its computation. This is easily implemented using MPI4Py through devices peer-to-peer communications. Relying on Tensorflow to perform said trainings, we encountered some unexpected issues.

The major one lies in the high variability in terms of VRAM consumption for each architecture probed. As direct consequence, upon not being able to estimate the consumption of the biggest architecture possible, one will be subject to Out Of Memory (OOM) errors that will perturb the search. However, it is not trivial to evaluate the precise maximum VRAM consumption for the search, even when you can roughly estimate the structure of your most complex network – this is mainly linked to the graph abstraction used by Tensorflow to encapsulate all part of the training process. Big providers of cloud DL platform have identified this issue and are currently studying it [22].

The second one is of a more nitpicking nature, but if not anticipated, one may incur a major code refactoring. In any Tensorflow backed training framework, the most commonly found import is the generic `import tensorflow as tf`. However, following this import, libraries opening are performed – e.g. CUDNN. In the context of a multiprocessing dispatching of training tasks on a multi-node multi-GPUs infrastructure, this may cause conflicts that terminate code execution. In a nutshell, for the specific case of dispatching previously explained,

this could very simply be avoided: the codebase should be structured in such a way that any `tensorflow import` is strictly performed at child process level. In this way, you can ensure that the task has been submitted properly – through MPI – to a dedicated GPU and can therefore proceed as any vanilla training perform on a single GPU computer.

Algorithm 1: MPI communication (discrete framework)

Result: Hist. storing arch. and corresponding performances
 $Orch. \leftarrow GPU_{intel};$
 $Orch. \xrightarrow[init.]{Pop., Hist.};$
while $Epoch$ **do**
 for $(i, arch.) \in enumerate(population - 1)$ **do**
 $Orch. \xrightarrow[arch]{GPU_i};$
 end
 $Train(arch)|_{Orch. \cup GPUs};$
 for $(i, arch.) \in enumerate(population - 1)$ **do**
 $Orch. \xleftarrow[arch, perf]{GPU_i};$
 $Orch. \xrightarrow[arch, perf]{Hist.};$
 end
 $Evolve(Pop.)|_{Orch.};$
end

A.2 Data parallelism: TensorFlow and Horovod

As training heavy search networks requires a significant amount of computing resources, we rely on data parallelism to keep the training time of such networks reasonable. In such a setting, two cases are to be considered:

- the usage of multiple GPUs on a single node, that we denote *mono-node/multi-gpu*
- the usage of multiple GPUs on multiple nodes, that we denote *multi-node/multi-gpu*

We use TensorFlow [35] as our reference deep learning framework, and so consider the native *mirrored strategy* available in the package. For our two use-cases, two distinct classes exist, respectively *MirroredStrategy* and *MultiWorkerMirroredStrategy*. Although we encountered no issue with the former, such a thing cannot be said for the latter.

We follow user guidelines and establish a generic codebase that handles both cases, based on *MultiWorkerMirroredStrategy*. As the majority of the methods we use are contained in an *experimental* package, and we are building a custom training loop, it seems reasonable to expect potential issues. As discussed before, while we encountered no difficulty working with a single node, it was not the case for multiple nodes. Most of the time, multi-node/multi-gpu trainings were stuck on a training or an evaluation step before reaching the end, most likely while applying reduction on all partial results from each worker, which implies a communication error between nodes. Such issues are also hard to solve as we do not get any feedback that can help identify the point of failure.

In a second step, we considered Horovod [36]. As we did not manage to solve the encountered TensorFlow issues, we relied on a dedicated and proven framework, widely used for distributed computing on HPC. Albeit this framework is harder to use initially – as any user needs to understand basic MPI concepts – it remains easy to integrate in an existing codebase. We updated our codebase to replace TensorFlow’s mirrored strategy by Horovod.

As the latter explicitly uses MPI concepts that are also used implicitly in the former, it is easier to understand and follow the training process. We encountered no blocking issue with Horovod for both use-cases, except some unusual behaviours (e.g deadlocks or crashes) that occurred only once. Additionally, switching to Horovod provided a small performance gain in terms of training time. Building upon these observations, we chose Horovod as our reference framework for data parallelism.