



HAL
open science

Deep semi-supervised approach based on consistency regularization and similarity learning for weeds classification

Farouq Benchallal, Adel Hafiane, Nicolas Ragot, Raphael Canals

► To cite this version:

Farouq Benchallal, Adel Hafiane, Nicolas Ragot, Raphael Canals. Deep semi-supervised approach based on consistency regularization and similarity learning for weeds classification. 2025. <hal-05403316>

HAL Id: hal-05403316

<https://hal.science/hal-05403316v1>

Preprint submitted on 8 Dec 2025

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

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



Distributed under a Creative Commons CC BY-NC-SA 4.0 - Attribution - Non-commercial use - ShareAlike - International License

DEEP SEMI-SUPERVISED APPROACH BASED ON CONSISTENCY REGULARIZATION AND SIMILARITY LEARNING FOR WEEDS CLASSIFICATION

A PREPRINT

✉ **Farouq Benchallal**

PRISME EA 4229
INSA CVL

Bourges, 18002, Centre Val de Loire, France
farouq.benchallal@insa-cvl.fr

✉ **Adel Hafiane**

PRISME EA 4229
INSA CVL

Bourges, 18002, Centre Val de Loire, France
adel.hafiane@insa-cvl.fr

✉ **Nicolas Ragot**

LIFAT EA 6300
Université Tours

Tours, 37200, Centre Val de Loire, France
nicolas.ragot@univ-tours.fr

✉ **Raphael Canals**

PRISME EA 4229
Université d'Orleans

Orleans, 45067, Centre Val de Loire, France
raphael.canals@univ-orleans.fr

October 14, 2025

ABSTRACT

Weed species classification represents an important step for the development of automated targeting systems that allow the adoption of precision agriculture practices. To reduce costs and yield losses caused by their presence. The identification of weeds is a challenging problem due to their shared similarities with crop plants and the variability related to the differences in terms of their types. Along with the variations in relation to changes in field conditions. Moreover, to fully benefit from deep learning-based methods, large fully annotated datasets are needed. This requires time intensive and laborious process for data labeling, which represents a limitation in agricultural applications. Hence, for the aim of improving the utilization of the unlabeled data, regarding conditions of scarcity in terms of the labeled data available during the learning phase and provide robust and high classification performance. We propose a deep semi-supervised approach, that combines consistency regularization with similarity learning. Through our developed deep auto-encoder architecture, experiments realized on the DeepWeeds dataset and inference in noisy conditions demonstrated the effectiveness and robustness of our method in comparison to state-of-the-art fully supervised deep learning models. Furthermore, we carried out ablation studies for an extended analysis of our proposed joint learning strategy.

Keywords Semi-supervised learning · Deep learning · Consistency regularization · Similarity learning · Precision Agriculture

1 Introduction

Weeds result in the highest potential yield loss in crops, along with pathogens (fungi, bacteria, etc.) and animal pests (insects, rodents, nematodes, mites, birds, etc.) being both of lesser importance [Oerke, 2006]. The losses attributed to weeds are related to several factors such as, their time of emergence, their density and type Chauhan [2020]. Weed species are highly competitive with crops for natural resources and characterized by a far better reproductive mechanism compared to crop plants. As a result, these unsown plant species constantly invade fields to overcome crops Kubiak et al. [2022], Radicetti and Mancinelli [2021]. The uncontrolled growth of weeds can lead to substantial economic

losses Chauhan [2020], Gharde et al. [2018], Llewellyn et al. [2016], Pimentel et al. [2005]. Therefore, addressing this problem is essential to meet the demands of an increasing global population. Currently, most weed management practices rely on the use of herbicides for the control of weeds. However, the over dependence on the application of herbicides with similar modes of action led to the development of herbicide-resistant weeds Heap [2025]. The growing concerns about the excessive use of agrochemicals Aktar et al. [2009] incentivize the adoption of precision agriculture practices Radoglou-Grammatikis et al. [2020].

The development of automated systems for weed control requires accurate identification and recognition of weed species. This is a challenging problem due to field conditions, such as changes in lighting and illumination and some cases the resemblance between weeds and crop plants regarding their color, texture and morphological characteristics [Hasan et al., 2021]. Employing machine learning approaches for weed species classification involves an extensive process of feature extraction and selection followed by the application of machine learning classifiers. Sabzi et al. Sabzi et al. [2020] extracted texture, color and shape features, along with five moment-invariant features from ground-based images. Subsequently, they applied different algorithms to select discriminative features, followed by weed classification using machine learning based classifiers, such as support vector machines and random forests. Machine learning methods require deep domain expertise to carry out time-intensive process of feature engineering. Conversely, deep learning approaches are distinguished by their strong ability to automatically extract discriminative features from data through representation learning Lecun et al. [2015]. The high learning capacity of the deep learning models enables them to perform classification and prediction remarkably well, given enough labeled data.

Currently, many successful deep neural network architectures Bah et al. [2018] have been used in agricultural applications. Reedha et al. Reedha et al. [2022] applied transformer-based networks, namely ViT-B-16 and ViT-B-32 for the classification of weeds and crop plants by means of unmanned aerial vehicles (UAVs). Ahmed et al. Ahmad et al. [2021] utilized convolutional neural networks (CNNs) such as ResNet50 He et al. [2016] and InceptionV3 Szegedy et al. [2015] for the identification of weed species within the fields of soybean and corn from ground-based imagery. Valente et al. Valente et al. [2022] applied seven variants of pre-trained CNN-models for the detection of weeds in grassland using UAV imagery. Based on cross validation method, the model MobileNet Howard et al. [2017] provided the best detection performance. Mesías-Ruiz et al. Mesías-Ruiz et al. [2024] utilized UAV images for weed species identification during the early growth stages in crop fields of maize and tomato, leveraging three types of CNNs i.e., VGG16 Simonyan and Zisserman [2015], ResNet152 He et al. [2016] and Inception-ResNet-V2 Szegedy et al. [2016]. In relation to the different training scenarios, Inception-ResNet-V2 achieved the highest performance.

The new developed saliency detection approaches Cong et al. [2019], Wang et al. [2022] favor deep learning-based techniques to discern salient objects and important regions in images. This is realized by relying on the multi-level features derived from deep neural networks to generate saliency maps. Zhao et al. Zhao and Wu [2019] proposed pyramid feature attention network to focus on high level context-aware features and on low-level spatial structural features, which were then both used to generate improved saliency maps. Furthermore, saliency detection can strengthen image classification and segmentation. Zeng et al. Zeng et al. [2019] developed a network that is composed of a segmentation model and saliency aggregation module to capture the connections between the two computer vision tasks. As for co-saliency detection, the goal is to detect salient regions from a set of related images by exploring the interdependencies between the images. Wei et al. Wei et al. [2017] presented a deep learning technique based on a fully convolutional network to detect co-salient objects. Similarly, co-segmentation aims to segment common objects from a group of relevant images. Li et al. Li et al. [2018] proposed a deep co-segmentation method to segment common objects belonging to a common semantic class from image pairs.

Nevertheless, to exploit the full potential of the deep learning algorithms, large, labeled datasets are needed, this represents a limitation for the agricultural applications as a consequence of the difficulties related to the annotation process. On the other hand, unlabeled data could be acquired in copious amounts with relative simplicity. The challenges to overcome regarding semi-supervised learning (SSL) is to perform similarly or surpass the state-of-the-art deep learning supervised models on datasets that are highly constrained in terms of labeled training data. A specific way to attain this is by solving the dilemma between leveraging an accurate model with enough parameters, considering that only a small amount of labeled data will be available for its training.

Our proposition is aligned with this direction. Henceforth, in this paper we propose a method that employs SSL to improve the way of utilizing the unlabeled data, in conditions where the labeled data is exceedingly more restricted. Our main contributions are outlined as follows:

- A deep semi-supervised method combining consistency regularization and similarity learning for effective and robust weed species recognition.

- An auto-encoder architecture, based on a ConvNeXt Encoder and specifically designed de coder with skip-connections. Through our joint learning strategy, it allows an increased ability for incorporating the most relevant information from unlabeled data in scenarios where the labeled data are excessively scarce.
- Extensive experiments on the public dataset DeepWeeds for a rigorous assessment and evaluation of our proposed approach to demonstrate its effectiveness and behavior.

The remainder of this paper is organized as follows. In the next section we present the related works. The subsequent section describes comprehensively our proposed method. The fourth section details at length the carried-out experiments alongside an in-depth analysis of the obtained results. Finally, our conclusions and perspectives are summarized in the last section.

2 Related Works

In the literature there are studies, which adopt the semi-supervised paradigm for the elaboration of methods with the objective of identifying and recognizing weed species. Kerdegari et al. Kerdegari et al. [2019] introduced a semi-supervised generative adversarial network (GAN) for the semantic segmentation of weeds and crops using aerial multispectral imagery. Khan et al. Khan et al. [2021] devised an optimized semi-supervised GAN-based framework for the identification of weeds and crops utilizing UAV images. Shorewala et al. Shorewala et al. [2021] applied a semi-supervised method which encompasses two steps. An unsupervised binary segmentation step to generate vegetation masks, followed by a classification step for distinguishing weeds and crops by means of a finetuned CNN model from ground-based imagery. Homan and du Preez. Homan and du Preez [2021] proposed a two fold method that is comprised of feature recognition along with the classification of plant species through deep semi-supervised learning. Hu et al. Hu et al. [2021] combined image synthesis with semi-supervised learning for the training of site-specific models to enable the detection of weeds. Liu et al. Liu et al. [2023] developed a semi-supervised approach that incorporates a mechanism of mixed attention to enhance the model’s ability in extracting important features to detect weed species. Kong et al. Kong et al. [2024] introduced the semi-supervised learning paradigm for the detection of weeds in wheat, showing that a full supervised method needs sufficient labeled data to achieve satisfactory performance compared to the SSL-based method. Similarly, Chen et al. Chen et al. [2024] used an SSL method for the detection of weeds in sod farms. Deng et al. Deng and Lu [2025] proposed a semi-supervised object detection method called WeedTeacher to benefit from the unlabeled data for enhanced detection of weeds in-domain and cross-domain contexts.

As aforementioned we developed our semi-supervised approach incorporating consistency regularization with similarity learning. With respect to the incorporation of consistency regularization we were mainly influenced by the ladder network Rasmus et al. [2015] due to its successful utilization of consistency regularization. Subsequently, several other consistency regularization methods were introduced, such as the Pi-model and temporal ensembling Laine and Aila [2016]. The Pi-model relies on stochastic transformations applied to the training samples for the integration of consistency regularization. Temporal ensembling relies on both applying stochastic transformations as well as leveraging an exponential moving average of the predictions. Following these methods the mean-teacher Tarvainen and Valpola [2017] was proposed, which relies more on the structure of the networks for the integration of consistency regularization. We should note that the methods that are based on consistency regularization suffer from confirmation bias Yang et al. [2021] due to the dependance on a single model within the semi-supervised architecture to generate the predictions for consistency training. In the case where these predictions (corresponding to the unlabeled training samples) are incorrect across multiple iterations, this would lead to a negative impact on the semi-supervised learning performance. To mitigate the effect of confirmation bias, Ke et al. Ke et al. [2019] presented an approach called dual student, which utilizes a stability constraint with respect to the unlabeled data to address the limitation of confirmation bias.

Relative to similarity learning we were motivated by methods that leverage joint learning for enhancing the predictive ability of the learning models. For instance, the deep learning framework introduced by Li et al. [2017], considers both class label information and local spatial distribution information between training samples through a pairwise loss to constrain representation learning. Similarly, Paclik et al. Paclik et al. [2006] proposed an approach which employs a trainable similarity measure to build second-stage classifiers for object detection. The measure is based on local matches in a set of regions within an image to increase robustness. Horn and Muller Horn and Müller [2018] proposed a neural network architecture termed similarity encoder constructed for simultaneously learning a mapping from an original input feature space into similarity preserving embedding space, while factorizing a target matrix with pairwise relations. Xia et al. Xia et al. [2015] developed a method that utilizes the cosine similarity metric for similarity learning, which is based on an ensemble of cosine learners.

Regarding our approach, we relied more on the design of our proposed deep encoder-decoder architecture and the utilization of skip-connections for the integration of consistency regularization for the purpose of limiting confirmation

bias by providing reliable information during the training from the labeled and unlabeled data. Moreover, to leverage our proposed architecture and develop our learning strategy, we were interested in how to further improve the consistency regularization training in relation to the scarcity of the labeled training data, by focusing on the incorporation of similarity learning into the training of the deep semi-supervised models to guide the process of joint learning toward improving the generalization ability of the trained deep semi-supervised models. Finally, the consistency-based methods mentioned above were applied to general purpose datasets. In contrast, for the development of our approach we were also motivated by the effective recognition and identification of weed species from images acquired in real world conditions.

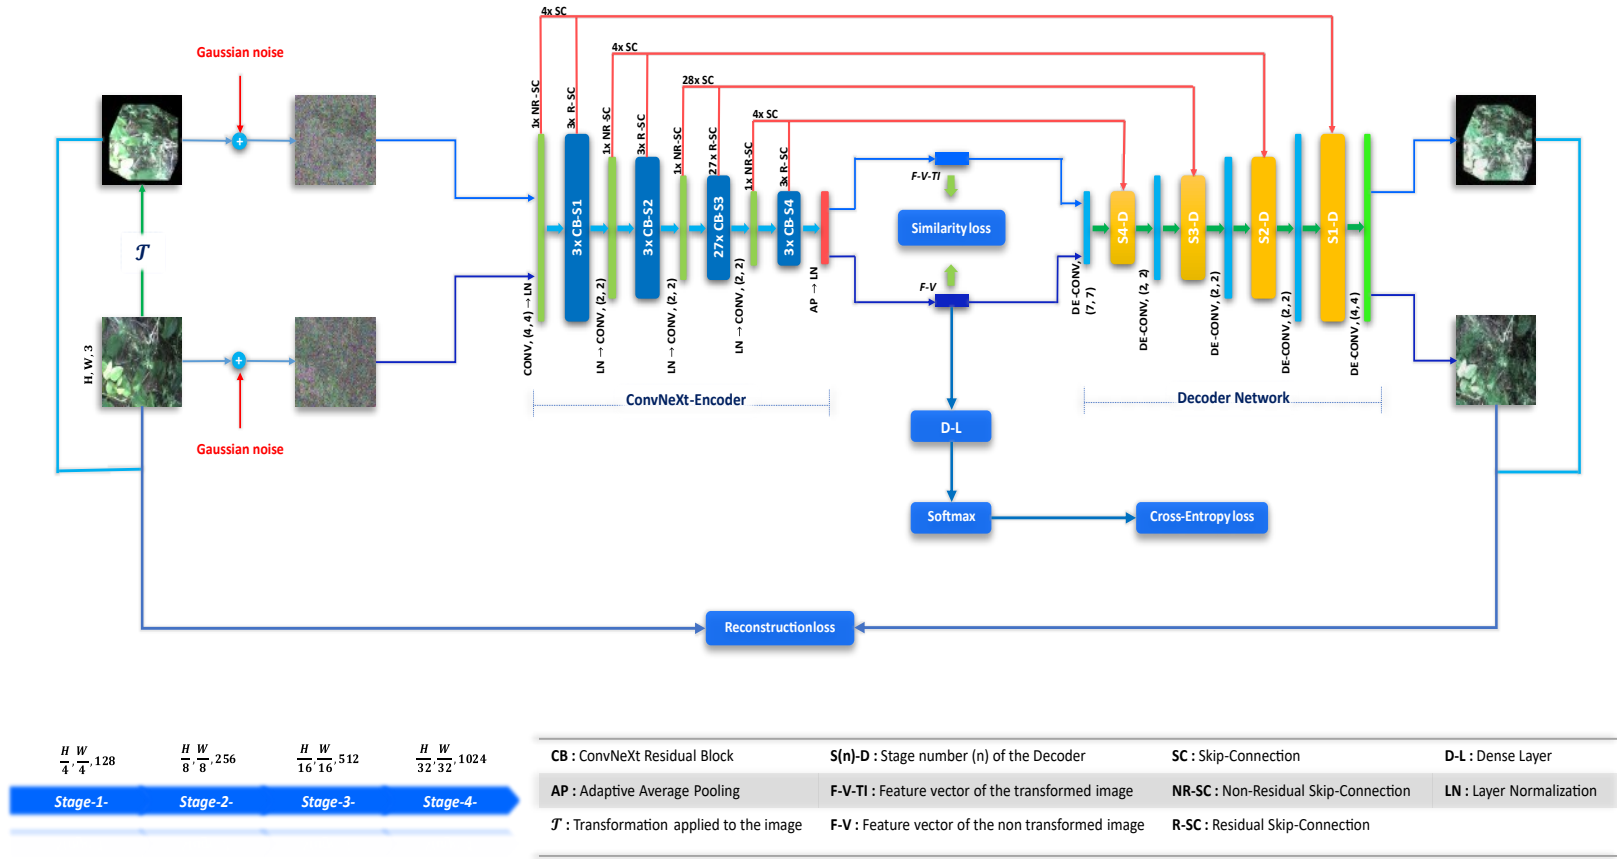


Figure 1: Overview of the semi-supervised approach

3 Proposed method

3.1 Problem statement

Supervised learning involves a sequence of training pairs $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)$, where each pair (\mathbf{x}_i, y_i) is composed of an input $\mathbf{x}_i \in \mathcal{X}$ from the space of inputs, and a label $y_i \in \mathcal{Y}$ from the space of outputs. In our case $\mathcal{X} \subset \mathbb{R}^d$ represents the set of images, and we have $\mathcal{Y} = \{1, \dots, C\}$ representing the species of weeds that we want to identify. C refers to the number of classes. The joint space $\mathcal{X} \times \mathcal{Y}$ is assumed to be a probabilistic space with an unknown probability measure $\mathcal{P}(\mathbf{x}, y)$ and the data is sampled from this space. The joint measure $\mathcal{P}(\mathbf{x}, y)$ can be decomposed into a measure of the marginal $\mathcal{P}(\mathbf{x})$, and a measure of the conditional distribution $\mathcal{P}(y|\mathbf{x})$. Supervised learning aims at estimating a functional relationship $\mathbf{x} \rightarrow y$ between a covariate $\mathbf{x} \in \mathcal{X}$ and the class variable $y \in \{1, \dots, C\}$, with the goal of minimizing the classification error. In semi-supervised learning settings, in addition to the labeled data $\mathcal{D}_\ell = \{(\mathbf{x}_i, y_i) | i = 1, \dots, n\}$ sampled from $\mathcal{P}(\mathbf{x}, y)$, we have also access to the unlabeled data $\mathcal{D}_u = \{\mathbf{x}_{n+j} | j = 1, \dots, m\}$ sampled from $\mathcal{P}(\mathbf{x})$. While acquiring labeled samples is expensive and time-intensive, gathering unlabeled samples can be done at low-cost and quickly. Then, $n \ll m$ which means that the size of the labeled portion could be much smaller than the size of the unlabeled portion Chapelle et al. [2006]. The objective of semi-supervised learning is to leverage the labeled data \mathcal{D}_ℓ , with additional information about the data distribution $\mathcal{P}(\mathbf{x})$, giving from \mathcal{D}_u with the purpose of increasing the performance and achieving better generalization to new unseen samples Oliver et al. [2018], Yang et al. [2021]. Semi-supervised learning requires that the data distribution should be under a set of assumptions Chapelle et al. [2006], Yang et al. [2021], otherwise, the prediction’s performance may not be improved. The main assumptions associated with semi-supervised learning are the following:

Semi-supervised smoothness assumption: if two points $\mathbf{x}_1, \mathbf{x}_2$ in a high-density region are close, then so should be the corresponding outputs y_1, y_2 .

Cluster assumption: If points are in the same cluster, they are likely to be of the same class, which means that the decision boundary should lie in a low-density region.

Manifold assumption: high-dimensional data lie approximately on a low-dimensional manifold. In high-dimensional spaces volume grows exponentially with the number of dimensions, which makes it difficult to estimate reliable densities. If the data lie on a low dimensional manifold, then the learning algorithms can overcome the problems related to high dimensionality, by operating in the corresponding low-dimensional space.

3.2 Architecture of the proposed Auto-Encoder

In the current section we describe in-depth both the encoder and decoder models utilized in our proposed deep auto-encoder architecture, depicted in Figure 1. This is followed by a detailed explanation of our semi-supervised learning strategy. Finally, we provide more details regarding our method of integrating the similarity loss, for the joint optimization process.

3.2.1 ConvNeXt-Encoder

ConvNeXt represents a family of pure convolutional neural networks devised through the gradual process of modernizing the standard family of RESNETS He et al. [2016] toward the design of vision transformers, through an in-depth exploration of the design spaces and the possible reachable limits, by relying solely on the use of CONVNET modules. The modernization process groups a series of design decisions applied on the macro/micro level-architecture [Liu et al., 2022]. These series of modernization steps can be outlined in the following manner. The first step involves the distribution of computation throughout the network, meaning the adoption of a stage compute ratio that is similar to the design of vision transformers Dosovitskiy et al. [2020], which specifies the number of blocks that will be attributed to each stage. Within a multi-stage design, the feature maps resolution is progressively changed across stages. The second step is related to the stem cell design, which is used for processing the input images at the start of the network. Following the design of hierarchical vision transformers, the RESNET style-stem cell is replaced with a patchify stem implemented through a 4×4 non-overlapping convolution layer. In the third step a special case of grouped convolution is adopted where the number of groups is equal to the number of channels. Depth-wise convolution Howard et al. [2017], Chollet [2016], Xie et al. [2016] bears resemblance to the weighted sum operation in self-attention, which mixes information only along the spatial dimension. Combining depth-wise convolution and 1×1 convolution results in a separation between the mixing of information along the spatial and the channel dimensions, which is an important property of vision transformers. Within the fourth step, the design of inverted bottleneck with an expansion rate equal to 4 is used. The decision was motivated by the design of the transformer block [Dosovitskiy et al., 2020, Vaswani et al., 2017] where an inverted bottleneck is created due to the hidden dimension of the MLP (MultiLayer Perceptron) block which is four times wider than the input dimension. In the fifth step, to benefit from large size kernels, the depth-wise convolution layer was positioned at the beginning of the residual block, followed with 1×1 convolution layer. This design choice is also inspired by the vision transformer block, in which we observe that the MSA (Multi-Head Self-Attention) module is positioned before the MLP layers. The size of the kernels associated with the depth-wise convolutional layer was increased from 3×3 to 7×7 . Regarding the sixth step, the Gaussian Error Linear Unit (GELU) Hendrycks and Gimpel [2016], which is the activation function utilized in most advanced transformers is used to replace the activation function Rectified Linear Unit (ReLU) Nair and Hinton [2010]. Subsequently, in order to replicate the style of the transformer block, all GELU activation functions within the residual block were removed with the exception of a single GELU activation function placed between two 1×1 convolution layers. Likewise, the number of batch normalization (BN) layers Ioffe and Szegedy [2015] was also reduced to a single layer per block. Moreover, the batch normalization layer was substituted with Layer normalization (LN) Ba et al. [2016] following the normalization method applied in vision transformers. As for the last step, to conduct the spatial downsampling between stages, separate downsampling layers were utilized instead of relying on the residual block at the start of each stage. And before each downsampling layer a LN layer was added. Regarding the patchify stem the normalization layer is added afterwards. The models of the ConvNeXt family achieve high-level performance across different model capacities (different types/variants of models) showing the effectiveness of the architecture design. The variant model that we have chosen for the encoder part of our proposed auto-encoder is called ConvNeXt-Base. We should note that the number of channels (the number of filters on the level of convolution layers) increases along the depth of the encoder as we move throughout the stages of the network.

3.2.2 Decoder network

Regarding the decoder network, we have adopted a multi-stage design with the aim of utilizing the feature maps with varying resolutions, from the different stages of the ConvNeXt encoder, for the reconstruction process. The feature maps on the encoder side are retrieved through residual and non-residual skip-connections. The residual skip-connections provide the decoder with direct access to the outputs of the residual blocks at any given stage of the encoder. The non-residual skip-connections enable the decoder to directly access the outputs of the downsampling layers including the patchify stem. The total number of skip-connections integrated into our architecture amounts to 40 (36 are residual skip-connections, corresponding to the ConvNeXt residual blocks and 4 non-residual skip-connections, corresponding to the downsampling layers and the patchify stem). We choose to utilize these two types of skip-connections to extract detailed information from the encoder path, which allows us to integrate the intermediate hidden representations (the encoder feature maps) in a manner that accounts for symmetry between the encoder and decoder, along with the dimensional change within the decoder path during the reconstruction process. Consistently with the design of

ConvNeXt-Base encoder, we have employed separate up sampling layers at the beginning of each decoder stage. These layers are essential for the structure of the decoder and for the process of integrating information from the skip-connections at each stage. Apart from the up-sampling layers, each decoder stage includes a specific number of de-convolution blocks. The number of these blocks corresponds to the number of skip-connections that connect the stage to the encoder path. Each de-convolution block receives as inputs the encoder representations from its respective skip-connection and the output of the preceding de-convolution block or the separate up-sampling layer. The source of the second input is determined by the position of the de-convolution block within the decoder stage. The first and the second inputs are then merged through element-wise summation, resulting in a mixing of the spatial information from both the decoder and the encoder. Subsequent to this step, a 1×1 de-convolution layer is applied with a stride of 1 and a number of kernels equal to the number of channels of the summed inputs. Due to the element-wise summation, the dimensions of the representations do not change within the de-convolution block. Thereafter, the output of the 1×1 de-convolution layer is normalized through layer normalization, and an activation function is applied in the final step to achieve the overall output of the block. With regard to the activation functions applied element-wise on the side of the decoder, we adopted LeakyReLU Maas [2013] for stages 4, 3 and 2. For stage 1, we applied the ELU activation function Clevert et al. [2016] that has improved learning characteristics and contributes to alleviate the vanishing gradient problem. As concluding element of the decoder, we added a final 4×4 de-convolution layer with a stride of 4 and 3 Kernels. An element-wise sigmoid activation function was then applied to generate the final output image. As part of preprocessing the pixel values of the input images were scaled between 0 and 1. Through dividing the channels of the images by 255. For semi-supervised learning settings a dense layer with a number of units equal to the classes that need to be recognized is added. Which is applied to the feature vector obtained by leveraging the encoder model.

3.3 Semi-supervised learning strategy

The semi-supervised optimization process is performed by minimizing the combined losses of both forms of data labeled and unlabeled. These combined losses comprise the following three distinct loss functions: supervised loss, consistency regularization loss (reconstruction from noisy inputs) and similarity loss. Their joint minimization using the available training data facilitates a refined guided optimization of the auto-encoder’s weights. The supervised loss is applied solely to the labeled portion of the training data, meaning that it relies on the supervised information provided by class labels of the annotated samples. Its minimization optimizes both the encoder model’s weights (θ_E) and the densely connected layer (θ_{d-l}) dedicated to classification. The cross-entropy function, defined below, is used to represent the supervised loss:

$$l_{CE}(y_i, f(\mathbf{x}_i + \zeta_i; \theta_E, \theta_{d-l})) = - \sum_{j=0}^{C-1} y_{ij} \log(f_j(\mathbf{x}_i + \zeta_i; \theta_E, \theta_{d-l})) \quad (1)$$

Here, y_i denotes the ground truth label. $f(\mathbf{x}_i + \zeta_i; \theta_E, \theta_{d-l})$ represents the predicted label \tilde{y}_i from the densely connected layer that employs the latent space from the encoder. ζ_i denotes the additive gaussian noise and C is the number of classes.

The consistency regularization loss utilizes both types of data labeled and unlabeled. As we mentioned previously, through the minimization of this loss we seek to exploit the characteristics of our proposed architecture that enables a significant flow of information during the backward pass. To compute this loss, we have used the L2 loss function and its formulation taking into consideration both the non-transformed view of the input image (\mathbf{x}_i) and its transformed view (\mathbf{x}_i') in the following way:

$$l_{CR}(\mathbf{x}_i, \tilde{\mathbf{x}}_i, \mathbf{x}_i', \tilde{\mathbf{x}}_i') = (1/2)(\|\mathbf{x}_i - \tilde{\mathbf{x}}_i\|_2^2 + \|\mathbf{x}_i' - \tilde{\mathbf{x}}_i'\|_2^2) \quad (2)$$

$\tilde{\mathbf{x}}_i$ denotes the reconstructed image of the non-transformed view and $\tilde{\mathbf{x}}_i'$ refers to the output reconstruction of the transformed view.

As for the similarity loss, both types of data are utilized too. Regarding the measure of similarity between the high-level representations, obtained through the encoder model, we used a loss which relies on the cosine similarity function (**cos**) defined as follows:

$$\begin{aligned} & \cos(f(\mathbf{x}_i + \zeta_i; \theta_E), f(\mathbf{x}_i' + \zeta_i'; \theta_E)) \\ &= f(\mathbf{x}_i + \zeta_i; \theta_E) \cdot f(\mathbf{x}_i' + \zeta_i'; \theta_E) / \max(\|f(\mathbf{x}_i + \zeta_i; \theta_E)\| \cdot \|f(\mathbf{x}_i' + \zeta_i'; \theta_E)\|, \epsilon) \end{aligned} \quad (3)$$

With $f(\mathbf{x}_i + \zeta_i; \theta_E)$ and $f(\mathbf{x}_i' + \zeta_i'; \theta_E)$ representing the feature vectors (high-level representations) corresponding to the non-transformed input image and its transformed version, respectively. ϵ is a constant set to 10^{-8} for avoiding division by zero. Henceforth, the similarity loss is expressed in the following way:

$$l_{Sim}(f(\mathbf{x}_i + \zeta_i; \theta_E), f(\mathbf{x}_i' + \zeta_i'; \theta_E)) = 1 - \cos(f(\mathbf{x}_i + \zeta_i; \theta_E), f(\mathbf{x}_i' + \zeta_i'; \theta_E)) \quad (4)$$

Finally, the expression of the overall loss composed of the three loss functions is given in the following manner:

$$l_{total} = l_{CE} + \lambda_{CR} \cdot l_{CR} + \lambda_{Sim} \cdot l_{Sim} \quad (5)$$

Where λ_{CR} and λ_{Sim} represent the two hyperparameters $\in [0, 1]$ that determine the relative contribution of the two losses to the overall loss.

3.4 Reconstruction from noisy images and similarity learning

As mentioned previously we have opted to combine the process of reconstructing clean outputs from noisy inputs using the auto-encoder, through consistency regularization constraint with similarity learning to increase further the ability to exploit unlabeled data. To enable this, the level where to incorporate the similarity term is important for the overall semi-supervised learning process.

Hence, our choice, with respect to where we want to employ this term was motivated by our interest in the extracted high-level abstract representations through the training data. These high-level representations of original image and transformed image are distinct because they contain high-level context aware information that are rich semantically. Thereby, we have chosen to integrate the similarity term at the end of the encoder network to guide the optimization process toward learning the most important characteristics shared between high-level representations. Furthermore, one should note that the transformations applied to generate the transformed views of the training samples do have an influence on the learning process. For this reason, these transformations need to be selected thoroughly to help in acquiring beneficial information for the process of similarity optimization (see Figure 2).




Figure 2: Resulting images reconstructed from noisy inputs through our proposed deep encoder-decoder model with similarity learning

4 Experiments and Results

This section introduces the experimental protocol and key evaluation metrics used for a thorough assessment of the proposed semi-supervised approach for weed species classification. The performance of our method is compared to state-of-the-art deep learning models that are characterized by having similar scale of trainable parameters, to highlight the benefits of the semi-supervised paradigm.

4.1 DeepWeeds dataset

For the purpose of training and evaluating the learning algorithms we have used the DeepWeeds dataset Olsen et al. [2019], which is constructed from images of weed species. Data were collected in-situ from eight different locations in northern Australia over a period spanning from June 2017 to March 2018. Eight weed species were selected (target classes) that have notoriety for their invasiveness and their damaging impact. In order to reflect scene and target variability, several factors of variations needed to be taken into account during the collection of these images: illumination, that varies throughout the day depending on the change in sunlight, which creates highly dynamic scenes; variability caused by complex and dynamic target backgrounds.

Furthermore, the dataset considers seasonal variation in the target weed species, which implies that a single class of weed species comprises images of the weed with and without flowers and in varying health conditions. These realistic environmental conditions constitute a significant challenge in the accurate classification of weed species. Additionally, DeepWeeds includes images of plant life native to the eight locations that represent negative class. Seeking to ensure the dataset diversity and breadth, more than 1,000 images per target weed species were gathered, with the objective of maintaining an approximate equal division of positive to negative class images across each location. In the end a total of 8,403 positive images of weed species and 9,106 negative images of neighboring flora and non-target backgrounds were collected. All the images within the DeepWeeds dataset were labeled by experts for the presence or absence of target weed species.

4.2 Deep learning models for comparative analysis

The supervised model **ConvNeXt-Base** Liu et al. [2022] is used as a basis for our encoder and is also used as a competitive approach. We have also incorporated two additional model variations from two separate neural network families: **EfficientNet-V2-L** Tan and Le [2021] and **ViT-B-16** Dosovitskiy et al. [2020]. These models were selected based on their performances and the availability of pre-trained weights on the extensive **ImageNet-22K** dataset, a superset of **ImageNet-1K** comprising 14,197,122 images categorized into $21,841 \approx 22K$ classes. Leveraging these pre-trained weights as a foundation for training on our target dataset offers considerable advantages, including improved performance and reduced training time compared to random weight initialization. **EfficientNet-V2-L** belongs to the **EfficientNet-V2** family, based on combining the training-aware neural architecture search and scaling for an optimized training speed and parameter efficiency. **ViT-B-16** on the other hand is part of the vision transformers family.

4.3 Settings for training and evaluation

4.3.1 Hyperparameter setup and implementation details

With respect to the implementation of the deep learning algorithms and the running of the experiments, we have used the machine learning framework PyTorch **1.10.2** as well as the libraries PyTorch-Ignite **0.4.8** and timm **0.5.4**. As for hardware specifications, we utilized the computing cluster CaSciModOT CaSciModOT [2025], which provided us access to three Nvidia Tesla V100 32 GB GPUs along with the AMD 7302 CPU.

To ensure a fair comparison of the results obtained by the deep learning models, we aimed to use consistent optimal hyperparameters across all experiments. Therefore, to update model parameters, we have used the stochastic gradient descent algorithm (**SGD**) with momentum and damping set to zero. As for learning rate, we have used two different values *i.e.*, 0.01 and 0.009. The learning rate is then decreased by a factor of 0.9 after each 10 epochs. Regarding the utilization of pre-trained weights, we have replaced the dense layer of ImageNet-22K with a layer that has a number of units equal to the number of classes of DeepWeeds (**9** units). In relation to our semi-supervised training approach, we have used a gaussian noise with a standard deviation σ to 0.1 and a mean μ set to zero. We have utilized pre-trained weights on the dataset ImageNet-1k, to initialize the decoders. We used two values for the λ_{sim} hyperparameter in our loss function: 0.90 and 0.85. The λ_{CR} hyperparameter was set to 1.00. To generate more unlabeled data, we have also applied rotations varying from -180° to $+180^\circ$ to the training subset. The obtained images were then de-labeled during the semi-supervised training.

4.3.2 Similarity transformations

As we noted previously, the transformations employed in the context of similarity learning are selected with attention to bring a wider range of diversity to the generated views from both labeled and unlabeled training images. The transformations comes with a series of geometric and intensity transformations applied throughout the course of the training epochs. Geometric transformations encompass rotations in the range of -120 and $+120$ degrees. A random shift, both horizontally (by a factor of 0.2) and vertically (by a factor of 0.3) with a probability of 0.7 is also applied. Additionally, a random scaling ranging from 0.8 to 0.9 with a probability of 0.7 was performed. Finally,

random horizontal and vertical flips were applied, each with a probability of 0.5. Intensity transformations consists in saturation transformations, which affect the intensity of colors, together with brightness transformations that change the illumination intensity. For the random sampling of the saturation factor, we have used two different values to represent the lower limit (min) of the sampling interval. The upper limit (max) was set to 2.6. In a similar manner we used the same low limit values for the random sampling of the brightness factor and the value 2.8 for the upper limit. The attributed probabilities for the saturation and brightness transformations are 0.6 and 0.7, respectively.

4.3.3 Evaluation procedure

During our experiments, we have employed the 5-fold stratified cross-validation method Arlot and Celisse [2010], Raschka [2018], Varma and Simon [2006]. Its aim is to determine the model’s ability to provide accurate predictions on new unobserved data samples. Stratified sampling was performed to ensure that each fold approximately reflects the original dataset’s class distribution. The initial partitioning of the dataset was done as follows: **60%** were allocated for training, **20%** used for validation and the remaining **20%** for testing, utilizing five rounds of cross-validation. Consequently, for each model type, five models in total are trained on various partitions of the data. We have used the validation subset for the selection of the optimal hyperparameters.

To meet our objectives of improving efficiency and predictive performance under constraints of limited labeled data, we reduced the labeled training samples to **20%** of the full dataset. Thereafter, we proceeded with a progressive decrease of the size of the labeled training subset by a factor of **5%** for each series of experiments. The lowest size of the labeled training portion used is equal to **5%** of the full dataset. Regarding the semi-supervised training, the unused labeled training images from the training set are de-labeled and used as part of the unlabeled training set. From this unlabeled set some will be selected based on a ratio between labeled and unlabeled training samples.

Moreover, to assess the impact of increasingly scarce labeled information on the performance of the models (semi-supervised and supervised), no data augmentation techniques were employed throughout the learning phase to expand the size of the labeled training subset. All models were trained for a duration of 60 epochs and for each model type within a specific family, the five best-performing models on the validation subset were selected and kept, due to the five rounds of cross-validation. Subsequently, the evaluation of these models on the test subset provides the final average performance discussed in this study.

The performance of the deep learning models was evaluated using two metrics namely: F1-Score and Accuracy Opitz and Burst [2019], Sokolova et al. [2006], Tharwat [2018]. The first metric, defined as the harmonic mean of two metrics: **1.** Precision, which indicates the proportion of correctly classified positive predictions relative to the total positive predictions ($FP + TP$). **2.** Recall denoting the proportion of correctly classified positive predictions relative to the real number of positives ($FN + TP$). For the second metric it is expressed as the ratio of correctly classified predictions to the total number of predictions ($FN + FP + TP + TN$). The equations defining these metrics are outlined as follows:

$$F1\text{-Score} = 2 \times ((Precision \times Recall) / (Precision + Recall)) \quad (6)$$

$$Precision = (TP) / (TP + FP), Recall = (TP) / (TP + FN) \quad (7)$$

$$Accuracy = (TP + TN) / (TP + TN + FP + FN) \quad (8)$$

4.4 Results of classification on the DeepWeeds dataset

Figures (3) and (4) depict the results of evaluations regarding the series of experiments conducted on the DeepWeeds dataset for both the semi-supervised and supervised models, based on the two metrics Accuracy and F1-Score. **ConvNeXt-Base-SSL-SCR** refers to the models trained in a semi-supervised manner using our auto-encoder architecture and the proposed learning strategy, which relies on a three terms of loss function: supervised loss, consistency regularization, and similarity loss. On the other hand, **ConvNeXt-Base-SSL** refers to the semi-supervised models trained in a similar way, through the proposed Encoder-Decoder structure incorporating two terms (supervised and consistency regularization) during the learning process, without applying the transformations of similarity.

From the obtained results, we can observe that when the size of the labeled training subset is constrained to **20%**, the semi-supervised models (ConvNeXt-Base-SSL-SCR and ConvNeXt-Base-SSL) achieved the highest classification performances (**92.51%** and **92.33%**) compared with the best performing supervised models (ViT-B-16 and ConvNeXt-

Base) with considerable margins in terms of both metrics. For instance, when we compare ConvNeXt-Base-SSL-SCR with ViT-B-16, we have a difference of **1.41%** in Accuracy and **2.24%** in F1-Score.

Moreover, when the size of the labeled training subset becomes more constrained moving from **20%** to **10%**, we can observe that the semi-supervised models continue to demonstrate higher performance scores (**89.23%** and **88.39%**) compared to the best performing supervised models, with more increased margins regarding Accuracy and F1-Score. When the size of the labeled training subset becomes very low, i.e., **5%**, we can see that ConvNeXt-Base-SSL underperforms in comparison to ConvNeXt-Base and ViT-B-16. In contrast, ConvNeXt-Base-SSL-SCR remains the best approach (**84.53%**) with more important margins given the limitations on the labeled training data. Similarly, when we carry out the comparison with respect to Efficientnet-V2-L, larger profound differences of **9.89%** in Accuracy and **15.1%** in F1-Score. The results of classification achieved by ConvNeXt-Base-SSL-SCR with a decreased labeled training subset of **5%**, show the importance of having greater degree of unlabeled data utilization, to enable a more refined ability of integrating information from the unlabeled training samples.

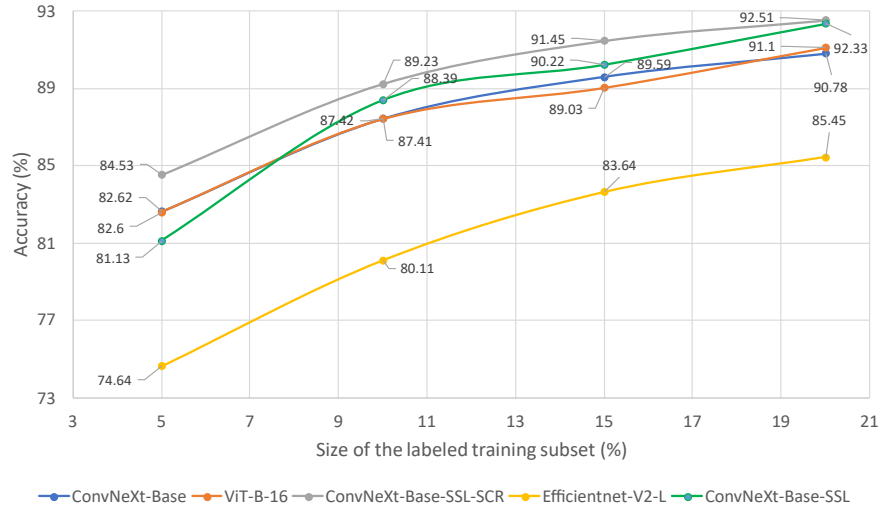


Figure 3: Average performance of the deep learning models (semi-supervised and supervised) on the test subset, based on the Accuracy metric, with different sizes of the labeled training subset. For **20%** labeled training subset, the results of the supervised models and ConvNeXt-Base-SSL were reported from Benchallal et al. [2024]

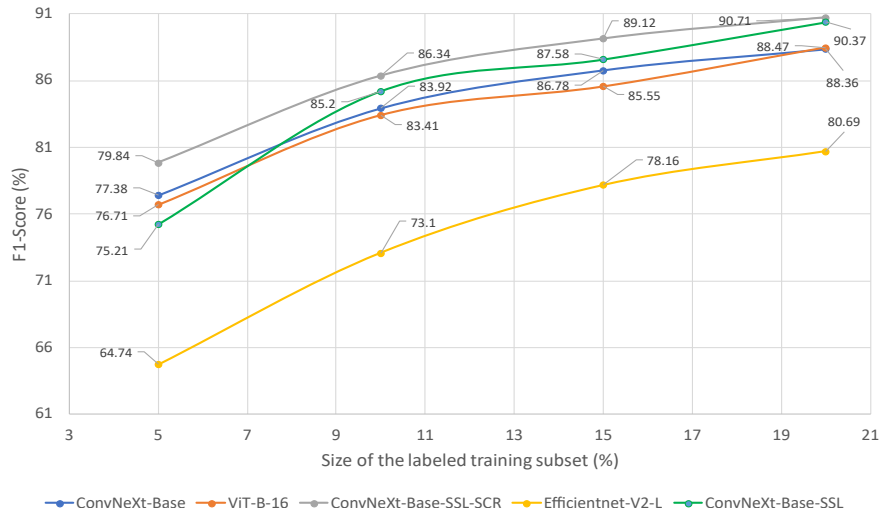


Figure 4: Average performance of the deep learning models (semi-supervised and supervised) on the test subset, based on the F1-Score metric, with different sizes of the labeled training subset

4.5 Noise influence with few labeled data




Figure 5: Average performance(Accuracy and F1-Score) of the deep learning models on the test subset, with added gaussian noise to the inputs. Data partitioning: **15%** labeled data, **20%** validation, **20%** test




Figure 6: Average performance (Accuracy and F1-Score) of the deep learning models on the test subset, with added gaussian noise to the inputs. Data partitioning: **10%** labeled data, **20%** validation, **20%** test




Figure 7: Average performance (Accuracy and F1-Score) of the deep learning models on the test subset, with added gaussian noise to the inputs. Data partitioning: **5%** labeled data, **20%** validation, **20%** test

Subsequent to the previous experiments, we have conducted an additional series of evaluations with the best-performing models saved from previous steps. The goal is to ascertain to which extent the presence of noise during inference combined with the scarcity of labeled data at training time affect the classification performance of both the semi-supervised and supervised models. These evaluations were undertaken using noisy images from the test subset, created through the addition of gaussian noise sampled from a normal distribution characterized by a fixed mean μ and a varying standard deviation σ . The classification results in this context of noisy conditions, using the models trained with **15%** and **10%**, respectively of labeled data, are illustrated in Figures 5 and 6. From these results we can observe that applying additive gaussian noise during the inference phase has a negative impact on the classification performance of purely supervised models. The impact becomes significant with raising values of the standard deviation σ . For example, in the case where the size of the labeled training subset is limited to **15%**, the performance of ViT-B-16 in terms of Accuracy drops from **89.03%** to **82.56%**, when the standard deviation σ is increased to **0.125**. As for the F1-Score metric, its performance drops from **85.55%** to **76.54%**. In a similar manner, a smaller proportion of labeled training data leads to a further drop in performance of the supervised models. Particularly, as we can observe in Figure 6, with an increased value of σ , the performance of ViT-B-16 drops to an Accuracy of **78.18%** and an F1-Score of **68.84%**. Figure 7 depicts the results of inference in noisy conditions corresponding to the models trained using only **5%** of

labeled training samples. Similarly to the previous cases, we can identify a consistent negative correlation between the importance of noise and the drop in performance of the supervised models. Overall, the drop in performance of **Efficientnet-V2-L** is much more profound compared to **ViT-B-16** and **ConvNeXt-Base**, regarding the effect of higher levels of noise and the limited availability of labeled training data. Conversely, the semi-supervised models demonstrated a robust and strong classification performance relative to the increase of noise importance, which can be observed across all the different scenarios under the limitation of labeled training data. Additionally, when considering the various raising values of the standard deviation σ , **ConvNeXt-Base-SSL-SCR** outperforms **ConvNeXt-Base-SSL** with important margins, especially in the case of less labeled training data.

4.6 Ablation studies




Figure 8: Comparative results of the average performance achieved by the semi-supervised models on the test subset, with and without the application of the similarity transformations (TFsim) in the loss. A ratio of one to five (1:5) between labeled to unlabeled data was adopted regarding both cases with respect to the size of the labeled training portion (15% and 20%)

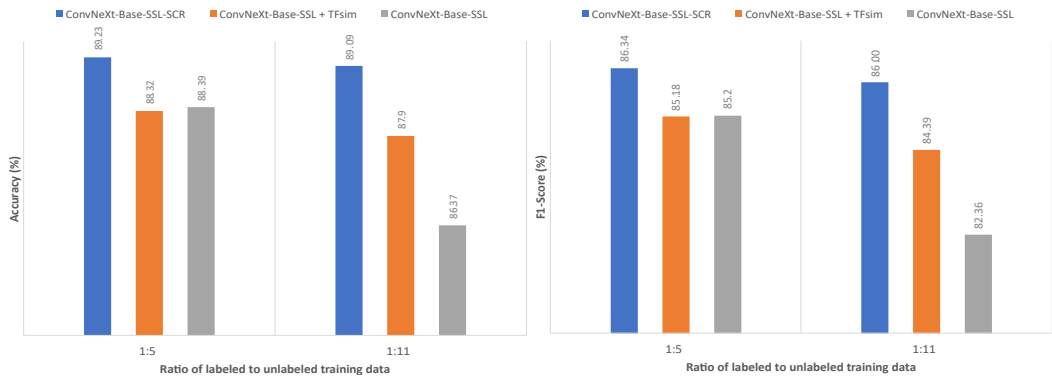


Figure 9: Comparative results of the average performance achieved by the semi-supervised models on the test subset, with and without applying the transformations of similarity in the loss with a portion of labeled training data equal to 10%

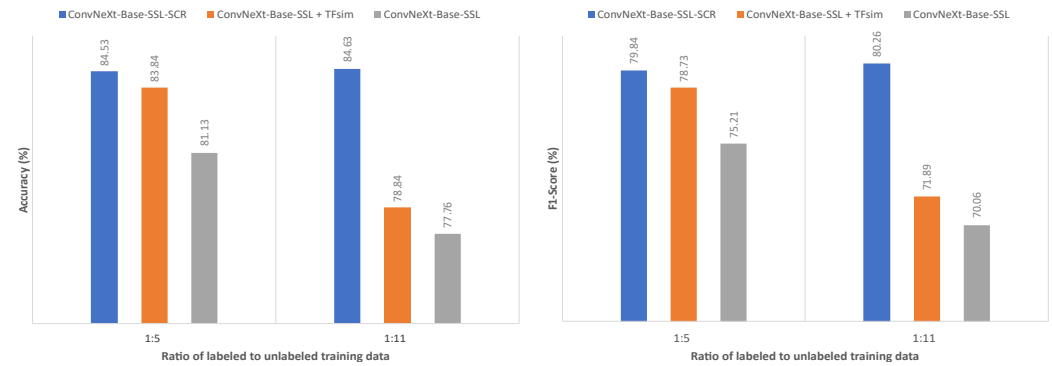


Figure 10: Comparative results of the average performance achieved by the semi-supervised models on the test subset, with and without applying the transformations of similarity in the loss with a portion of labeled training data equal to 5%

In these experiments we realized a set of ablations focusing specifically on our semi-supervised models. Here, we have adopted the same optimal hyperparameters employed as before. This experimental setup aims to provide a more in-depth assessment of our proposed strategy, focusing on the influence of the labeled-to-unlabeled data ratio and the impact of incorporating similarity transformations into the loss function. Figure 8 illustrates the results obtained by the semi-supervised models corresponding to the cases where the size of the labeled training subset is constrained to **20%** and **15%**, respectively, with a ratio between labeled to unlabeled data of **1:5**. When the labeled training subset is limited to **15%**, we can observe that utilizing two terms of the total loss along with the transformations of similarity (**ConvNeXt-Base-SSL + TFsim**) yields improvement over the performance of the semi-supervised models (**ConvNeXt-Base-SSL**) trained using two terms without applying the transformations of similarity. For **20%** of labeled data, this conclusion is less noticeable. With respect to a limitation of **10%** of the labeled training data, shown in Figure 9, with two different ratios utilized between labeled and unlabeled data. We can see that the impact of similarity loss depends on the ratio between labeled and unlabeled data. With a ratio of one to eleven (**1:11**), the semi-supervised models utilizing the transformations of similarity provide a better accuracy compared to the ConvNeXt-Base-SSL models with only two terms in loss. However, with a ratio of one to five (**1:5**) between the labeled and unlabeled data both variants of the semi-supervised models (ConvNeXt-Base-SSL TFsim and ConvNeXt-Base-SSL) achieved approximately similar classification results (**88.39%** in Accuracy and **85.20%** in F1-Score). On the other hand, the deep semi-supervised models (ConvNeXt-Base-SSL-SCR) trained through the utilization of all three combined terms of the total loss achieved higher classification performance in comparison to the other two types of the semi-supervised models with respect to both applied ratios between labeled to unlabeled data. Similarly, for the case where the labeled training portion is restricted to **5%**, we can observe from Figure 10, that the ConvNeXt-Base-SSL-SCR models continue to outperform the other variants of the semi-supervised models, especially when the ratio between labeled to unlabeled data is higher. These results demonstrate the effectiveness of our proposed joint learning strategy regarding the ability to incorporate pertinent and relevant information from the unlabeled training data.

5 Conclusions

In this paper we introduced a deep semi-supervised approach which combines consistency regularization training with similarity learning, by leveraging our proposed auto-encoder architecture. The proposed strategy of joint learning aims at further improving the ability of incorporating information from unlabeled data, when the labeled training data become increasingly scarce. This approach was motivated by the need to accurately classify weed species from images captured in real-world conditions. To this end, we rely on deep semi-supervised models trained using both labeled and unlabeled data. The realized series of experiments on the dataset DeepWeeds showed that the models trained by integrating all three terms of our joint learning strategy provided higher classification results, particularly under conditions of limited labeled training data, outperforming both fully supervised models and semi-supervised models trained leveraging only the supervised and consistency regularization terms of the overall loss function. Furthermore, inference in noisy conditions with few labeled data demonstrated the robustness of our method compared to traditional models. Moreover, we conducted extended ablation experiments to assess the influence of the ratio between labeled to unlabeled data and the importance of the transformations applied regarding similarity learning. In future work we intend to further develop the joint learning strategy through the incorporation of pseudo labeling and extend the application of our method to object detection.

Declarations

- **Funding**

This work was carried out as a part of DESHERBROB project funded by Region Centre-Val de Loire, France. We gratefully acknowledge its support.

- **Conflict of interest/Competing interests**

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**

We have shared the link to the public dataset DeepWeeds used to carry out the experiments.

- **Authors contributions**

Farouq Benchallal: Conceptualization, Methodology, Software, Visualization, Writing– original draft, Investigation, Validation. **Adel Hafiane:** Conceptualization, Methodology, Writing– review & editing, Supervision, Project administration, Funding acquisition, Validation. **Nicolas Ragot:** Conceptualization, Methodology, Writing– review & editing, Supervision, Validation. **Raphael Canals:** Conceptualization, Methodology, Writing– review & editing, Supervision, Project administration, Validation.

References

- E. C. Oerke. Crop losses to pests. *Journal of Agricultural Science*, 144:31–43, 2006. ISSN 00218596. doi:10.1017/S0021859605005708.
- Bhagirath Singh Chauhan. Grand challenges in weed management. *Frontiers in Agronomy*, 1, 2020. ISSN 26733218. doi:10.3389/fagro.2019.00003.
- Adrianna Kubiak, Agnieszka Wolna-Maruwka, Alicja Niewiadomska, and Agnieszka A. Pilarska. The problem of weed infestation of agricultural plantations vs. the assumptions of the european biodiversity strategy. *Agronomy*, 12(8), 2022. ISSN 2073-4395. doi:10.3390/agronomy12081808. URL <https://www.mdpi.com/2073-4395/12/8/1808>.
- Emanuele Radicetti and Roberto Mancinelli. Sustainable weed control in the agro-ecosystems. *Sustainability (Switzerland)*, 13, 2021. ISSN 20711050. doi:10.3390/su13158639.
- Yogita Gharde, P. K. Singh, R. P. Dubey, and P. K. Gupta. Assessment of yield and economic losses in agriculture due to weeds in india. *Crop Protection*, 107:12–18, 2018. ISSN 02612194. doi:10.1016/j.cropro.2018.01.007.
- Rick Llewellyn, David Ronning, Michael Clarke, Allan Mayfield, Allan Mayfield Consulting, and Steve Walker. *Impact of Weeds on Australian Grain Production : The cost of weeds to Australian grain growers and the adoption of weed management and tillage practices. Report for Grains Research and Development Corporation*. 2016. ISBN 9781921779916. URL www.grdc.com.au/bookshop.
- David Pimentel, Rodolfo Zuniga, and Doug Morrison. Update on the environmental and economic costs associated with alien-invasive species in the united states. *Ecological Economics*, 52:273–288, 2005. ISSN 09218009. doi:10.1016/j.ecolecon.2004.10.002.
- Ian Heap. Current status of the international herbicide-resistant weed database. <https://www.weedscience.org>, 2025. Accessed on 2025-05-02.
- Wasim Aktar, Dwaipayan Sengupta, and Ashim Chowdhury. Impact of pesticides use in agriculture: Their benefits and hazards. *Interdisciplinary Toxicology*, 2:1–12, 2009. ISSN 13379569. doi:10.2478/v10102-009-0001-7.
- Panagiotis Radoglou-Grammatikis, Panagiotis Sarigiannidis, Thomas Lagkas, and Ioannis Moscholios. A compilation of uav applications for precision agriculture. *Computer Networks*, 172, 2020. ISSN 13891286. doi:10.1016/j.comnet.2020.107148.
- A. S.M.Mahmudul Hasan, Ferdous Sohel, Dean Diepeveen, Hamid Laga, and Michael G.K. Jones. A survey of deep learning techniques for weed detection from images. *Computers and Electronics in Agriculture*, 184, 2021. ISSN 01681699. doi:10.1016/j.compag.2021.106067.
- Sajad Sabzi, Yousef Abbaspour-Gilandeh, and Juan Ignacio Arribas. An automatic visible-range video weed detection, segmentation and classification prototype in potato field. *Heliyon*, 6, 5 2020. ISSN 24058440. doi:10.1016/j.heliyon.2020.e03685.
- Yann Lecun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521:436–444, 2015. ISSN 14764687. doi:10.1038/nature14539.
- M Dian Bah, Adel Hafiane, and Raphael Canals. Deep learning with unsupervised data labeling for weed detection in line crops in uav images. *Remote Sensing*, 10(11), 2018. ISSN 2072-4292. doi:10.3390/rs10111690. URL <https://www.mdpi.com/2072-4292/10/11/1690>.
- Reenul Reedha, Eric Dericquebourg, Raphael Canals, and Adel Hafiane. Transformer neural network for weed and crop classification of high resolution uav images. *Remote Sensing*, 14(3), 2022. ISSN 2072-4292. doi:10.3390/rs14030592. URL <https://www.mdpi.com/2072-4292/14/3/592>.
- Anis Ahmad, Dharmendra Saraswat, Varun Aggarwal, Aaron Etienne, and Benjamin Hancock. Performance of deep learning models for classifying and detecting common weeds in corn and soybean production systems. *Computers and Electronics in Agriculture*, 184:106081, 2021. ISSN 0168-1699. doi:<https://doi.org/10.1016/j.compag.2021.106081>. URL <https://www.sciencedirect.com/science/article/pii/S0168169921000995>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi:10.1109/CVPR.2016.90.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. 2015. URL <http://arxiv.org/abs/1512.00567>.
- João Valente, Santosh Hiremath, Mar Ariza-Sentís, Marty Doldersum, and Lammert Kooistra. Mapping of rumex obtusifolius in nature conservation areas using very high resolution uav imagery and deep learning. *International Journal of Applied Earth Observation and Geoinformation*, 112:102864, 2022. ISSN 1569-8432. doi:<https://doi.org/10.1016/j.jag.2022.102864>. URL <https://www.sciencedirect.com/science/article/pii/S1569843222000668>.

- Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017. URL <https://arxiv.org/abs/1704.04861>.
- G.A. Mesías-Ruiz, I. Borra-Serrano, J.M. Peña, A.I. de Castro, C. Fernández-Quintanilla, and J. Dorado. Weed species classification with uav imagery and standard cnn models: Assessing the frontiers of training and inference phases. *Crop Protection*, 182:106721, 2024. ISSN 0261-2194. doi:<https://doi.org/10.1016/j.cropro.2024.106721>. URL <https://www.sciencedirect.com/science/article/pii/S0261219424001492>.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015. URL <https://arxiv.org/abs/1409.1556>.
- Christian Szegedy, Sergey Ioffe, and Vincent Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *CoRR*, abs/1602.07261, 2016. URL <http://arxiv.org/abs/1602.07261>.
- Runmin Cong, Jianjun Lei, Huazhu Fu, Ming-Ming Cheng, Weisi Lin, and Qingming Huang. Review of visual saliency detection with comprehensive information. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(10): 2941–2959, 2019. doi:10.1109/TCSVT.2018.2870832.
- Wenguan Wang, Qiuxia Lai, Huazhu Fu, Jianbing Shen, Haibin Ling, and Ruigang Yang. Salient object detection in the deep learning era: An in-depth survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6): 3239–3259, 2022. doi:10.1109/TPAMI.2021.3051099.
- Ting Zhao and Xiangqian Wu. Pyramid feature selective network for saliency detection. *CoRR*, abs/1903.00179, 2019. URL <http://arxiv.org/abs/1903.00179>.
- Yu Zeng, Yun-Zhi Zhuge, Huchuan Lu, and Lihe Zhang. Joint learning of saliency detection and weakly supervised semantic segmentation. *CoRR*, abs/1909.04161, 2019. URL <http://arxiv.org/abs/1909.04161>.
- Lina Wei, Shanshan Zhao, Omar El Farouk Bourahla, Xi Li, and Fei Wu. Group-wise deep co-saliency detection. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3041–3047, 2017. doi:10.24963/ijcai.2017/424. URL <https://doi.org/10.24963/ijcai.2017/424>.
- Weihao Li, Omid Hosseini Jafari, and Carsten Rother. Deep object co-segmentation. *CoRR*, abs/1804.06423, 2018. URL <http://arxiv.org/abs/1804.06423>.
- Hamideh Kerdegari, Manzoor Razaak, Vasileios Argyriou, and Paolo Remagnino. Semi-supervised GAN for classification of multispectral imagery acquired by uavs. *CoRR*, abs/1905.10920, 2019. URL <http://arxiv.org/abs/1905.10920>.
- Shahbaz Khan, Muhammad Tufail, Muhammad Tahir Khan, Zubair Ahmad Khan, Javaid Iqbal, and Mansoor Alam. A novel semi-supervised framework for uav based crop/weed classification. *PLoS ONE*, 16, 2021. ISSN 19326203. doi:10.1371/journal.pone.0251008.
- Shantam Shorewala, Armaan Ashfaq, R. Sidharth, and Ujjwal Verma. Weed density and distribution estimation for precision agriculture using semi-supervised learning. *IEEE Access*, 9:27971–27986, 2021. ISSN 21693536. doi:10.1109/ACCESS.2021.3057912.
- Dewald Homan and Johan A. du Preez. Automated feature-specific tree species identification from natural images using deep semi-supervised learning. *Ecological Informatics*, 66:101475, 2021. ISSN 1574-9541. doi:<https://doi.org/10.1016/j.ecoinf.2021.101475>. URL <https://www.sciencedirect.com/science/article/pii/S1574954121002661>.
- Chengsong Hu, J. Alex Thomasson, and Muthukumar V. Bagavathiannan. A powerful image synthesis and semi-supervised learning pipeline for site-specific weed detection. *Computers and Electronics in Agriculture*, 190:106423, 2021. ISSN 0168-1699. doi:<https://doi.org/10.1016/j.compag.2021.106423>. URL <https://www.sciencedirect.com/science/article/pii/S0168169921004403>.
- Teng Liu, Xiaojun Jin, Luyao Zhang, Jie Wang, Yong Chen, Chengsong Hu, and Jialin Yu. Semi-supervised learning and attention mechanism for weed detection in wheat. *Crop Protection*, 174:106389, 2023. ISSN 0261-2194. doi:<https://doi.org/10.1016/j.cropro.2023.106389>. URL <https://www.sciencedirect.com/science/article/pii/S0261219423002120>.
- Xiaotong Kong, Teng Liu, Xin Chen, Peng Lian, Danlan Zhai, Aimin Li, and Jialin Yu. Exploring the semi-supervised learning for weed detection in wheat. *Crop Protection*, 184:106823, 2024. ISSN 0261-2194. doi:<https://doi.org/10.1016/j.cropro.2024.106823>. URL <https://www.sciencedirect.com/science/article/pii/S0261219424002515>.
- Xin Chen, Teng Liu, Kang Han, Xiaojun Jin, and Jialin Yu. Semi-supervised learning for detection of sedges in sod farms. *Crop Protection*, 179:106626, 2024. ISSN 0261-2194. doi:<https://doi.org/10.1016/j.cropro.2024.106626>. URL <https://www.sciencedirect.com/science/article/pii/S0261219424000541>.

- Boyang Deng and Yuzhen Lu. Semi-supervised weed detection in vegetable fields: In-domain and cross-domain experiments, 2025. URL <https://arxiv.org/abs/2502.17673>.
- Antti Rasmus, Harri Valpola, Mikko Honkala, Mathias Berglund, and Tapani Raiko. Semi-supervised learning with ladder networks. 2015. URL <http://arxiv.org/abs/1507.02672>.
- Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. 2016. URL <http://arxiv.org/abs/1610.02242>.
- Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. 2017. URL <http://arxiv.org/abs/1703.01780>.
- Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu. A survey on deep semi-supervised learning. *CoRR*, abs/2103.00550, 2021. URL <https://arxiv.org/abs/2103.00550>.
- Zhanghan Ke, Daoye Wang, Qiong Yan, Jimmy S. J. Ren, and Rynson W. H. Lau. Dual student: Breaking the limits of the teacher in semi-supervised learning. *CoRR*, abs/1909.01804, 2019. URL <http://arxiv.org/abs/1909.01804>.
- Ya Li, Xinmei Tian, Xu Shen, and Dacheng Tao. Classification and representation joint learning via deep networks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 2215–2221, 2017. doi:10.24963/ijcai.2017/308. URL <https://doi.org/10.24963/ijcai.2017/308>.
- P. Paclik, J. Novovicova, and R.P.W. Duin. A trainable similarity measure for image classification. In *18th International Conference on Pattern Recognition (ICPR'06)*, volume 3, pages 391–394, 2006. doi:10.1109/ICPR.2006.188.
- F. Horn and K.-R. Müller. Predicting pairwise relations with neural similarity encoders. *Bulletin of the Polish Academy of Sciences Technical Sciences*, 66(No 6 (Special Section on Deep Learning: Theory and Practice)):821–830, 2018. doi:10.24425/bpas.2018.125929. URL <http://journals.pan.pl/Content/109871/PDF-MASTER/DOI.pdf>.
- Peipei Xia, Li Zhang, and Fanzhang Li. Learning similarity with cosine similarity ensemble. *Information Sciences*, 307: 39–52, 2015. ISSN 0020-0255. doi:<https://doi.org/10.1016/j.ins.2015.02.024>. URL <https://www.sciencedirect.com/science/article/pii/S0020025515001243>.
- Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien, editors. *Semi-Supervised Learning*. The MIT Press, 2006. ISBN 9780262033589. doi:10.7551/mitpress/9780262033589.001.0001. URL <https://doi.org/10.7551/mitpress/9780262033589.001.0001>.
- Avital Oliver, Augustus Odena, Colin Raffel, Ekin D. Cubuk, and Ian J. Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. *CoRR*, abs/1804.09170, 2018. URL <http://arxiv.org/abs/1804.09170>.
- Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CoRR*, abs/2201.03545, 2022. URL <https://arxiv.org/abs/2201.03545>.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *CoRR*, abs/2010.11929, 2020. URL <https://arxiv.org/abs/2010.11929>.
- François Chollet. Xception: Deep learning with depthwise separable convolutions. *CoRR*, abs/1610.02357, 2016. URL <http://arxiv.org/abs/1610.02357>.
- Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. *CoRR*, abs/1611.05431, 2016. URL <http://arxiv.org/abs/1611.05431>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017. URL <http://arxiv.org/abs/1706.03762>.
- Dan Hendrycks and Kevin Gimpel. Bridging nonlinearities and stochastic regularizers with gaussian error linear units. *CoRR*, abs/1606.08415, 2016. URL <http://arxiv.org/abs/1606.08415>.
- Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In Johannes Fürnkranz and Thorsten Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 807–814. Omnipress, 2010. URL <https://icml.cc/Conferences/2010/papers/432.pdf>.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. URL <http://arxiv.org/abs/1502.03167>.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016. URL <https://arxiv.org/abs/1607.06450>.
- Andrew L. Maas. Rectifier nonlinearities improve neural network acoustic models. 2013. URL <https://api.semanticscholar.org/CorpusID:16489696>.

- Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1511.07289>.
- Alex Olsen, Dmitry A. Konovalov, Bronson Philippa, Peter Ridd, Jake C. Wood, Jamie Johns, Wesley Banks, Benjamin Girgenti, Owen Kenny, James Whinney, Brendan Calvert, Mostafa Rahimi Azghadi, and Ronald D. White. Deepweeds: A multiclass weed species image dataset for deep learning. *Scientific Reports*, 9, 2019. ISSN 20452322. doi:10.1038/s41598-018-38343-3.
- Mingxing Tan and Quoc V. Le. Efficientnetv2: Smaller models and faster training. *CoRR*, abs/2104.00298, 2021. URL <https://arxiv.org/abs/2104.00298>.
- CaSciModOT. Regional computing cluster cascimodot. <https://cascimodot.fr/>, 2025. Accessed on 2025-05-04.
- Sylvain Arlot and Alain Celisse. A survey of cross-validation procedures for model selection. *Statistics Surveys*, 4: 40–79, 2010. ISSN 19357516. doi:10.1214/09-SS054.
- Sebastian Raschka. Model evaluation, model selection, and algorithm selection in machine learning. *CoRR*, abs/1811.12808, 2018. URL <http://arxiv.org/abs/1811.12808>.
- Sudhir Varma and Richard Simon. Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*, 7, 2006. ISSN 14712105. doi:10.1186/1471-2105-7-91.
- Juri Opitz and Sebastian Burst. Macro F1 and macro F1. *CoRR*, abs/1911.03347, 2019. URL <http://arxiv.org/abs/1911.03347>.
- Marina Sokolova, Nathalie Japkowicz, and Stan Szpakowicz. Beyond accuracy, f-score and roc: A family of discriminant measures for performance evaluation. volume WS-06-06, pages 24–29, 2006. ISBN 1577352882. doi:10.1007/11941439_114.
- Alaa Tharwat. Classification assessment methods. *Applied Computing and Informatics*, 17:168–192, 2018. ISSN 22108327. doi:10.1016/j.aci.2018.08.003.
- Farouq Benchallal, Adel Hafiane, Nicolas Ragot, and Raphaél Canals. Convnext based semi-supervised approach with consistency regularization for weeds classification. *Expert Systems with Applications*, 239:122222, 2024. ISSN 0957-4174. doi:10.1016/j.eswa.2023.122222. URL <https://www.sciencedirect.com/science/article/pii/S0957417423027240>.