

ResNeTS: a ResNet for Time Series Analysis of Sentinel-2 Data Applied to Grassland Plant-Biodiversity Prediction

Álvaro G. Dieste, Francisco Argüello, Dora B. Heras, *Member, IEEE*, Paul Magdon, Anja Linstädter, Olena Dubovyk, and Javier Muro

Abstract—Analyzing time series from remote sensing data can aid in understanding spectral-temporal phenomena in ecosystems, such as the seasonal variation of plant components. Lately, deep learning has emerged as a strong method for mapping environmental variables from this data due to its exceptional predictive capabilities. This work studies the adaptation of the ResNet computer vision architecture for time series analysis of Sentinel-2 data. The resulting deep learning architecture, ResNeTS, stacks sequential convolutions to build a deep and narrow network, aligning with the design principles of leading convolutional architectures in computer vision. Experiments were carried out for predicting different plant-biodiversity indices, namely species richness, and Shannon and Simpson indices, for temperate grassland ecosystems. The results show that ResNeTS can achieve moderate improvements in terms of accuracy compared to other state-of-the-art architectures, such as InceptionTime (up to +0.021 r^2), with reduced computational costs owing to its streamlined architecture.

Index Terms—biodiversity prediction, deep learning, multispec-

tral imaging, remote sensing, residual network, sentinel-2, time series analysis

ABBREVIATIONS

CNN Convolutional Neural Network
FCN Fully Convolutional Network
Grad-CAM Gradient-weighted Class Activation Mapping
GRU Gated Recurrent Unit
LAI Leaf Area Index
LSTM Long Short-Term Memory
MLP Multilayer Perceptron
MODIS Moderate Resolution Imaging Spectroradiometer
NDVI Normalized Difference Vegetation Index
 r^2 Coefficient of Determination
ResNet Residual Network
ResNeTS Residual Network for Time Series
RMSE Root Mean Squared Error
rRMSE Relative Root Mean Squared Error
sRMSE Systematic Root Mean Squared Error
uRMSE Unsystematic Root Mean Squared Error
RNN Recurrent Neural Network
SAR Synthetic Aperture Radar
SGD Stochastic Gradient Descent
SHAP Shapley Additive Explanations

This work was supported in part by the German Research Foundation (DFG) under the Priority Program 1374, through the Project Sensing Biodiversity Across Scales (SeBAS, Grants DU 1596/1-1 and LI 1842/4-1). This work was also supported in part by grants PID2022-141623NB-I00 and TED2021-130367B-I00 funded by MCIN/AEI/10.13039/501100011033 and by “European Union NextGenerationEU/PRTR”; by Xunta de Galicia - Consellería de Cultura, Educación, Formación Profesional e Universidades [Galician Research Center accreditation ED431G-2019/04 and Reference Competitive Group accreditation ED431C-2022/16]; and by the European Regional Development Fund (ERDF/EU). The work of Álvaro G. Dieste was also supported in part by Xunta de Galicia - Consellería de Cultura, Educación, Formación Profesional e Universidades, under Grant ED481A-2022/257. (Corresponding author: Álvaro G. Dieste.)

Álvaro G. Dieste and Dora B. Heras are affiliated with the Singular Research Center on Intelligent Technologies (CiTIUS), University of Santiago de Compostela, 15782 Santiago de Compostela, Spain (e-mail: alvaro.goldar.dieste@usc.es; dora.blanco@usc.es).

Francisco Argüello is affiliated with the Department of Electronics and Computing, University of Santiago de Compostela, 15782 Santiago de Compostela, Spain (e-mail: francisco.arguello@usc.es).

Paul Magdon is affiliated with the University of Applied Sciences and Arts (HAWK), 37077 Göttingen, Germany (e-mail: paul.magdon@hawk.de).

Anja Linstädter is affiliated with the Biodiversity Research/Systematic Botany, University of Potsdam, 14469 Potsdam, Germany (e-mail: linstaedter@uni-potsdam.de).

Olena Dubovyk is affiliated with the Geography Institute, Department of Earth System Sciences, University of Hamburg, 20146 Hamburg, Germany (e-mail: olena.dubovyk@uni-hamburg.de).

Javier Muro is affiliated with the Institute of Farm Economics, Thünen Institute, 38116 Braunschweig, Germany (e-mail: javier.muro@thuenen.de).

The code developed for this work will be available at <https://github.com/alvrogd/ResNeTS>. All data necessary to replicate the study is available at the repository of the Biodiversity Exploratories BExIS (<https://www.bexis.uni-jen.a.de/>); accession numbers 30969 (Sentinel-2 time series) and 27386 (Species inventories).

I. INTRODUCTION

The current biodiversity crisis has tremendous impacts on the proper functioning of ecosystems and its derived benefits to humans [1]. For instance, biodiversity is vital for food production as it directly influences the maintenance of fundamental ecosystem functions, such as soil fertilization, pest and disease regulation, erosion control, biomass production, and pollination of crops and trees [2], [3]. In this context, biodiversity conservation, i.e., protecting and preserving the richness and variety of ecosystems, habitats, and species, is of imperative importance [4], [5]. This is particularly true for European grasslands, which are an integral part of Europe's agricultural landscapes [6]. They have a key role in ecosystem health and functionality, as they act as habitats for many plant and animal species. Human intervention, including the transformation of natural grasslands, intensive cultivation, pollution, and climate change, can affect the diversity and stability of grassland ecosystems in many ways and at different scales. In recent decades, the decline in grassland diversity across Europe has threatened biodiversity and is a major environmental concern [7].

To effectively address these challenges, the application of remote sensing data from satellite images has become

indispensable [8], [9], [10]. In particular, time series analysis of sequences of satellite images allows for the extraction of temporal characteristics critical for understanding ecosystem dynamics. By capturing seasonal variations with this data, researchers can enrich the understanding of ecosystems, helping the development of precise solutions for environmental challenges such as land cover classification [11] and phenological metrics extraction [12].

Building on this technological foundation, researchers have initiated multiple efforts to monitor grassland biodiversity, using temporal data extracted from satellite imagery at different scales. The biodiversity mapping includes alpha diversity, which quantifies the biodiversity within a community or ecosystem [13], [14]; beta diversity, which measures the diversity between two communities [15]; and gamma diversity, a combination of the other two metrics that assesses the biodiversity at landscape scale between different ecosystems [16]. Another way to map biodiversity is using radiative transfer models to map functional diversity: the different strategies for growth, reproduction, and survival of organisms that determine ecosystem services, resilience, and stability [17]. Much of this research is confined to hyperspectral sensors (e.g., plane or UAV mounted), with a much more limited spatial and temporal coverage than multispectral sensors (e.g., Landsat, Sentinel-2). This considerably reduces the practical applicability of these approaches for systematic monitoring [18]. Additionally, the problem of spatial autocorrelation in empirical models is too often not addressed, and thus, model results cannot be trusted outside the training area [19].

Deep learning approaches have recently shown high predictive capabilities and a unique ability to reveal higher hierarchical patterns from spectral-temporal data like time series [20]. The application of these powerful models can offer substantial advantages for biodiversity monitoring using remote sensing data, facilitating the move beyond traditional land cover type classifications and spurious correlations between biodiversity and spectral indices [21]. This approach also circumvents the reliance on controversial assumptions, such as the spectral diversity hypothesis [22]. While the applications of deep learning methods for biodiversity monitoring are relatively scarce, their potential for time series analysis of remote sensing data can be appreciated in a broader context of applications:

- 1) **Multilayer Perceptron (MLP)** [23]: This architecture marked a first milestone by outperforming traditional machine learning approaches in analyzing time series. Numerous studies have since explored and continued using it for a wide range of applications. For instance, [24] used an MLP to estimate biomass and species richness of grasslands in three regions of Germany using Sentinel-2 surface reflectance, achieving better estimations than a Random Forest regressor [25]. Other interesting applications of MLPs are described in [26] and [27]. Despite their capabilities, a notable drawback of MLPs is their limited ability to model temporal and dimensional patterns in multivariate time series effectively. This limitation arises from their simultaneous analysis of all data points, without considering the sequential nature of the data. To overcome these limitations, more advanced deep

learning architectures like the Recurrent Neural Network (RNN) [28] and the Convolutional Neural Network (CNN) [29] were eventually introduced.

- 2) **RNNs**: These include the Long Short-Term Memory (LSTM) [30] and Gated Recurrent Unit (GRU) [31]. They process time series data step by step, updating an internal state that models relations with previous steps. For instance, [32] used both a LSTM and a GRU to generate a global Leaf Area Index (LAI) from Moderate Resolution Imaging Spectroradiometer (MODIS) data with better accuracy compared to an MLP approach. More advanced applications use their bidirectional variants, Bi-LSTM and Bi-GRU, which process time series in both forward and backward directions to understand temporal patterns better. For example, [33] leveraged a Bi-LSTM architecture to detect rice crops in southern Brazil from Sentinel-1 Synthetic Aperture Radar (SAR) time series, [34] used another Bi-LSTM for crop classification from Sentinel-2 data, and [35] employed a Bi-GRU architecture to estimate Normalized Difference Vegetation Index (NDVI) in two areas in France using SAR and optical data from Sentinel-2. More innovative approaches incorporate additional components, such as attention mechanisms [36], into these bidirectional architectures to further enhance their capabilities [37], [38].
- 3) **CNNs**: While initially developed for visual information processing, they also exhibit potential for time series analysis. Originally, CNNs slide 2D convolutional kernels across an image, extracting dimensional and spatial patterns between pixels and their neighbors to capture visual features effectively. In the early attempts of using CNNs for time series analysis, a common approach involved converting the time series data into images through feature engineering [39], [40]. This transformation aimed to map the time series temporal and dimensional aspects onto the image representation's height and width, but it proved to be suboptimal compared to RNNs [39], [41]. Eventually, CNNs started to be specifically designed for time series analysis using 1D convolutions, thus outperforming RNNs in numerous remote sensing studies. For instance, [42] used a 1D CNN for early crop classification in the Zhanjiang region of China using Sentinel-1 SAR time series, outperforming both LSTM and GRU architectures. Another study by [43] showcased the advantages of a 1D CNN over a LSTM for crop classification in California, USA, using Landsat surface reflectance. Following the footsteps of computer vision, time series CNNs evolved from shallow architectures with few layers to deeper, more powerful architectures, commonly known as residual CNNs [44]. These networks can be recognized by the presence of skip-connections between layers, which allow for constructing architectures with tens or hundreds of layers while circumventing training issues such as gradient vanishing. In particular, [45] introduced the first residual CNN with widespread adoption in time series analysis, including remote sensing applications. For instance, [46] demonstrated the capabilities of this

architecture for mapping different irrigation systems across Catalonia, Spain, using data from multiple satellites.

- 4) **Transformers [36]:** This relatively new architecture has impacted a wide range of disciplines due to its excellent capabilities when compared to other deep learning architectures [47]. Transformers analyze time series by leveraging their novel self-attention mechanism, which simultaneously weighs the importance of all time steps while capturing complex dependencies across the entire sequence, without the sequential processing limitations of RNNs. Although the application of Transformers in analyzing time series from remote sensing data is still emerging, they have already been successfully employed in various tasks such as image land cover classification [48], [49].

Despite the recent attention on Transformers, CNNs have continued to demonstrate competitive performance in various areas, particularly in computer vision, as exemplified by architectures like ConvNeXt [50]. Moreover, CNNs benefit from the inherent inductive bias of convolutions, making them less data-intensive during training compared to Transformers [51]. This trait can be particularly helpful for remote sensing applications, where datasets often contain limited data [52].

Among the current CNN architectures for time series analysis, InceptionTime [53] stands out for its high accuracy [54] and widespread use across various domains, including remote sensing [55]. Inspired by the Inception-v4 network from image processing [56], InceptionTime consists of modules that perform parallel convolutions of different sizes, coupled with residual connections, to handle time series of diverse lengths.

However, most innovations in CNN development are driven by image processing, with other fields like time series analysis often experiencing a delayed adoption of these new ideas. Since the introduction of Inception-v4 in 2017, image processing CNNs have increasingly favored deep, narrow networks based on the Residual Network (ResNet)¹ family [44]. These models prioritize a higher number of sequential modules interconnected with residual connections [57], [58], [50]. Developing a ResNet-based architecture for time series analysis could accelerate progress in the area by facilitating the integration of cutting-edge innovations from the most advanced image processing CNNs.

All these observations have inspired this work to develop the Residual Network for Time Series (ResNeTS), a deep learning architecture based on the ResNet family, now adapted for time series analysis of remote sensing data. ResNeTS is evaluated on its ability to predict three biodiversity indices over grasslands with different species compositions, using Sentinel-2 time series. In particular, ResNeTS is benchmarked against InceptionTime and other state-of-the-art architectures for time series analysis, including Transformers. Unlike InceptionTime's branching topology with parallel convolutions, ResNeTS features a simpler architecture of sequential convolutions. This design approach, when considering the characteristics of the time

series to analyze, improves the accuracy of InceptionTime while reducing computational costs. These results illustrate the potential of continuing CNN development for precise and efficient time series analysis in remote sensing tasks.

The rest of this manuscript is organized as follows. The study area and the remote sensing dataset are presented in Section II. Then, the ResNeTS architecture for time series analysis of remote sensing data is detailed in Section III. Thereafter, the experiments for evaluating the efficacy of ResNeTS for the prediction of plant-biodiversity indices of interest are presented in Section IV; an explainability analysis is conducted as well to assess the importance of each spectral and temporal features used. Next, the implications of the results for biodiversity research and management, and for further time series analysis in a broader scenario of remote sensing applications, are discussed in Section V. Finally, the main conclusions of this work are summarized in Section VI.

II. MATERIALS

A. Study area

The Biodiversity Exploratories encompass three distinct regions distributed across Germany, Central Europe (Figure 1). The three regions are the UNESCO Biosphere Reserve "Schorfheide-Chorin" (SCH), the National Park "Hainich" and its surroundings (HAI), and the Biosphere Reserve "Schwäbische Alb" (ALB). Across and within these sites, species richness and compositions are highly variable, as they are influenced by abiotic environmental conditions such as topographic position, soil characteristics, and land use intensities [59], [60].

Each exploratory region comprises 50 experimental grassland plots, representing a spectrum of management practices ranging from intensive to extensive land use. Some plots are subjected to varying degrees of cattle grazing, while others are mown at specific seasonal intervals. Additionally, different levels of fertilization are implemented. To ensure standardized quantification of management levels, a Land Use Intensity (LUI) index was developed [59]. This index is calculated annually based on factors such as mowing frequency, grazing intensity, and fertilization inputs.

Figure 2 shows examples of different degrees of species diversity in seminatural grasslands across Germany. The grassland plots in the Biodiversity Exploratories are usually dominated by perennial grass species, such as meadow soft grass (*Holcus lanatus*), tall oat grass (*Arrhenatherum elatius*), or smooth meadow grass (*Poa pratensis*). Among the legume species, white clover (*Trifolium repens*), red clover (*Trifolium pratense*), and bird's-foot trefoil (*Lotus corniculatus*) are frequently found. As forbs, commonly occurring species are dandelion (*Taraxacum officinale*), wild carrot (*Daucus carota*), or plantain (*Plantago lanceolata*). Some of these and other species found at the Biodiversity Exploratories are shown in Figure 3.

Using biodiversity indices can help assess the health and stability of grassland ecosystems as well as their provided functions and services. This is also vital for grassland management and conservation. Among these indices, the following are particularly relevant and will be the object of study in this work:

¹The term "ResNet" will be used to refer to the family of residual networks introduced in [44], whereas "residual CNN" will be used to refer to derived networks (i.e., CNNs that integrate skip-connections in their structure).




Figure 1. (a) Location of the three study sites (Sch: Schorfheide-Chorin Biosphere Reserve, Hai: Hainich National Park, Alb: Schwäbische Alb) in Germany, Central Europe. Colours depict altitude (m.a.s.l.), based on the Shuttle Radar Topography Mission digital terrain model. (b) Illustration of a monitoring plot (50 × 50 m) and the reduced sampling area (4 × 4 m) where plant species inventories were collected, based on a false color UAV image. (c) Example of two adjacent grasslands with different land-use intensities, resulting in different compositions and diversity of plant communities.



Figure 2. Examples of different degrees of grassland biodiversity in the temperate regions of Germany. Biodiversity increases from left to right.

- *Species richness*: The species number per unit of area (S).
- *Shannon index*:

$$H = - \sum_{i=1}^S p_i \cdot \ln(p_i), \quad (1)$$

where H = Shannon's diversity index, S = Total number of species, and p_i = Relative abundance of species i .

- *Simpson index*:

$$D = 1 - \frac{\sum_{i=1}^S (n_i \cdot (n_i - 1))}{N \cdot (N - 1)}, \quad (2)$$

where D = Simpson's diversity index, n_i = Number of individuals of the species i , N = Total number of

individuals in the community, and S = Total number of species.

Simpson's index reflects both species richness (S) and evenness (the distribution of species relative abundances within a community). Here, 1 indicates infinite diversity (all species are equally present), and 0 indicates no diversity (only one species present) [61]. Shannon's index also considers both species richness and evenness, but is more sensitive to rare or less abundant species [61]. It ranges between 0 (low richness and evenness) and 5, but most values range between 1.5 and 3.5.




Figure 3. Examples of plant species present in the grasslands of the Biodiversity Exploratories, with three grass species shown in the top row. (a) *Alopecurus pratensis*. (b) *Phleum pratensis*. (c) *Arrhenatherum elatius*. (d) *Trifolium pratense* (a legume species). (e) *Daucus carota* (a forb species). (f) *Ranunculus bulbosus* (another forb species).

B. Remote sensing and species datasets

This study uses the floristic and remote sensing datasets compiled in [24]. Plant species inventories were obtained during the peak standing crop period (second half of May), from 2017 to 2020 (4 years), following a standardized sampling protocol across the 150 plots [62] (total of 600 measures). Such inventories were recorded in a permanently marked 4×4 meter area within each of the plots (Figure 1b), which is considered as representative of the 50×50 meter plot [60]. Vascular plant species' abundances were recorded annually via visual estimations of their ground cover. From this data, diversity indices (S , H , and D) were calculated for each plot. Plots with trees within or near their boundaries were excluded from the analysis, resulting in 502 valid samples. Spatial dependencies within site and across sites for this dataset were analyzed and proved non-significant in [24].

The Sentinel-2 images utilized in this study were processed using the FORCE software [63]. Images from 2017–2020 with a cloud cover lower than 50% were downloaded, atmospherically and topographically corrected, and resampled to 10 m. Subsequently, a synthetic time series was generated by creating an image every 14 days through radial basis function interpolation implemented in the Time Series Analysis module of the FORCE processor. Images acquired during winter (November–March) were excluded from further analysis due to their limited significance for vegetation studies and more

frequent cloud cover. The result is a synthetic gapless dataset of 10 bands and 16 time steps (see Figure 4) representative of the spectral variation of the vegetation across the phenological cycles of several years. The median pixel value for each band was calculated for each plot and paired with each field observation, as in [24].

The resulting dataset is composed of 502 observations of each biodiversity metric (either species richness S , Shannon's diversity index H , or Simpson's diversity index D), and their corresponding temporal profiles (March–October with 16 time steps) of reflectance values for 10 bands of Sentinel-2 [64]: blue, green, red, red edge 1, red edge 2, red edge 3, near-infrared, near-infrared b, short-wave infrared 1 and short-wave infrared 2. The coastal blue, water vapor, and cirrus bands were excluded.

Deep learning studies tend to train architectures with several thousands of observations acquired from different sources [65]. Hence, the reduced size of the dataset used in this study may seem strange, but it has various reasons. First, biodiversity observations for a single dataset must be manually collected in situ, using standardized sampling protocols such as those undertaken at the Biodiversity Exploratories, resulting in laborious annual field campaigns [66]. This inevitably limits the size of a single dataset. Second, it is often not feasible to compile local biodiversity (alpha-diversity) data from different sources for a joint analysis. For instance, species richness




Figure 4. Species-rich (left column) and species-poor (right column) plots along with their temporal profile in the near-infrared and short-wave infrared bands of Sentinel-2 for 2020.

is directly influenced by the sampling area: the larger the sampling area, the higher the number of recorded species [67]. Inventories acquired on different spatial scales are thus difficult to compare. Unfortunately, a standardization of species richness on sampling area is also not straightforward and usually requires elaborate methods to estimate the shape of the species-area relationship [67], [68], [69]. Finally, even though the Biodiversity Exploratories started in 2008, the Sentinel-2 B was not launched until 2017, further limiting the number of observations in the dataset.

III. RESNETS

A. ResNeTS overview

ResNeTS is a **ResNet**-based network for time series analysis of Sentinel-2 data. It favors a simple design, stacking sequential convolutions to construct a narrow network. By considering in its design the characteristics of time series to analyze (i.e., length and dimensionality), **ResNeTS** achieves state-of-the-art accuracy with reduced computational costs.

As any **ResNet**-based architecture, **ResNeTS**' structure is divided into three sections, as shown in the left side of Figure 5: stem, residual blocks, and head. The stem prepares the input time series for the residual blocks, which then extract




Figure 5. Overview of ResNeTS. (a) Three sections of the ResNeTS architecture. (b) Application of a 1D convolution to analyze a time series. The convolution uses a filter of 5 elements, producing an output series with more bands (taller) than the input.

meaningful features through the repeated use of convolutions. Finally, the head generates the predictions using these features.

Each convolution analyzes the time series sliding its filters step by step, aiming to extract temporal and dimensional patterns through the simultaneous analysis of the current step and its neighboring steps across all dimensions. This temporal window incorporating information from all dimensions is depicted on the right side of Figure 5.

Compared to shallow CNNs, the integration of shortcuts in residual networks allows the construction of significantly deeper architectures without encountering common training issues such as gradient vanishing. Thus, by layering a high number of convolutional layers and progressively increasing the number of filters, ResNeTS extracts increasingly complex patterns from the input data for enhanced prediction accuracy.

As previously discussed, InceptionTime, the most notable CNN for time series analysis, consists of a series of blocks that apply parallel convolutions to the data. This concept is illustrated on the left part of Figure 6, where the four parallel convolutions are applied and combined as input for the subsequent block, along with the skip connection (left convolution). In contrast, the building blocks of ResNeTS consist of just two sequential convolutions and the corresponding skip-connection (left path).

Comparing both architectures, fewer and simpler blocks are required in ResNeTS to build an architecture that achieves an accuracy improved over InceptionTime. This results in a

considerably more lightweight topology, as depicted in the right part of Figure 6. This is because each InceptionTime block uses four parallel convolutional filters to adapt to short, medium, and long time series, whereas ResNeTS' convolutions are tuned for the target time series.

B. Initial ResNeTS architecture

The ResNet family from computer vision [44] consists of architectures with varying sizes, denoted by the total number of layers: ResNet-18, 34, 50, 101, 152. Architecture selection for a specific problem depends on both the available computational resources and the dataset size, as larger networks can assimilate more knowledge. Given the small size of the dataset used in this study, particularly when compared to large-scale computer vision datasets like ImageNet [70], the ResNet-18 architecture with 2D convolutions replaced by 1D convolutions set the initial baseline for ResNeTS, as it provided results comparable in accuracy to larger variants.

After that, simplifications were introduced to further reduce the size of the architecture without incurring any accuracy loss, as ResNet-18 was still too large for the dataset of interest. In particular:

- The number of repetitions per residual block was reduced from the default of two to one.
- The number of convolutional channels was reduced to 64 in all blocks.




Figure 6. Comparison between InceptionTime and ResNeTS architectures. (a) Depiction of the convolutional blocks of both architectures. (b) Depiction of the depth of both architectures.

The optimal channel configuration was obtained from an extensive grid search on the training points of the dataset. Let B_x represent the channel count of block x . In ResNet-18, the channel count is: $B_1 = 64$, $B_2 = 128$, $B_3 = 256$, $B_4 = 512$. For ResNeTS, each block was allowed to go from 32 convolutional channels to its maximum (e.g., 128 for the second block), with a step of 32, and ensuring that: $B_1 \leq B_2 \leq B_3 \leq B_4$.

Lastly, to make the ResNeTS baseline suitable for the target time series, most of the strides inherited from ResNet-18 for the progressive dimensionality reduction had to be removed. This was necessary as the series consists of only 16 time steps, whereas ResNet was initially designed for 112×112 RGB images. The four original strides in ResNet result in a total dimensionality reduction of $\times 1/16$, transforming the 112×112 image sample into a 7×7 one that enters the head layers. By retaining only one stride in ResNeTS (see Table I), the input for the head has a final size of 8 time steps, closely mirroring the original dimensions in ResNet.

Figure 7 shows the network topology of the resulting baseline ResNeTS architecture, containing the four residual blocks detailed in Table I.

C. Larger convolutional kernels

The use of 1D convolutions in CNNs for time series instead of 2D convolutions leads to a significant reduction in computational costs. Thus, much larger convolutional filters than the standard 3×3 filters in ResNet can be used, enabling a more comprehensive understanding of temporal and dimensional patterns in time series while keeping the architecture computationally feasible. In particular, all of the convolutional layers in ResNeTS used filters of size 5, as Table I details, with no further accuracy improvements observed beyond this point for the dataset used in this study.

D. Architectural improvements

Three minor modifications were introduced to ResNeTS to improve its accuracy further in predicting the biodiversity indices of interest. The first change was an upgrade in the skip-connections, while the other modifications were two layer-wise modifications.

ResNeTS inherited its skip-connections from the ResNet family. These present a design limitation that was addressed in ResNeTS as discussed in [71]. Specifically, when traversing feature maps of different spatial sizes (i.e., downsampling the processed time series), the original skip-connections also applied a stride to its 1-size convolutions intended for channel mapping, thus disregarding part of the input feature map.

Table I
DETAILS OF THE LAYERS IN **RESNETS** ARCHITECTURE. THE OUTPUT SIZE CORRESPONDS TO AN INPUT TIME SERIES WITH DIMENSIONS $length \times bands = 16 \times 10$, MATCHING THE DATASET USED IN THIS STUDY.

Component	Layer	Output size	Filter size	Stride	Normalization	Activation
Stem	1D Conv	16×96	1	1	Batch Norm	ReLU
Residual Block	1D Conv	16×64	5	1	Batch Norm	ReLU
	1D Conv	16×64	5	1	Batch Norm	–
	1D Skip Conv	16×64	1	1	Batch Norm	–
	Addition	16×64	–	–	–	ReLU
Residual Block	1D Conv	16×64	5	1	Batch Norm	ReLU
	1D Conv	16×64	5	1	Batch Norm	–
	1D Skip Conv	16×64	1	1	Batch Norm	–
	Addition	16×64	–	–	–	ReLU
Residual Block	1D Conv	8×64	5	2	Batch Norm	ReLU
	1D Conv	8×64	5	1	Batch Norm	–
	1D Skip Average Pool	8×64	2	2	–	–
	1D Skip Conv	8×64	1	1	Batch Norm	–
	Addition	8×64	–	–	–	ReLU
Residual Block	1D Conv	8×64	5	1	Batch Norm	ReLU
	1D Conv	8×64	5	1	Batch Norm	–
	1D Skip Conv	8×64	1	1	Batch Norm	–
	Addition	8×64	–	–	–	ReLU
Head	1D Average Pool	1×64	8	8	–	–
	Flatten	64	–	–	–	–
	Fully Connected	1	–	–	–	Softmax

This approach is shown on the left side of Figure 8. By incorporating an extra average pooling layer within these connections when needed, as depicted on the right side of Figure 8, all feature points are utilized during downsampling. This enables the subsequent convolutional layer to concentrate solely on mapping the channel count to produce the skip-connection. Note how the final **ResNeTS** architecture, depicted in Figure 7, included this pooling operation in the third block, which performed the spatial downsampling from 16 time steps to 8.

Second, since not all residual blocks in **ResNeTS** have a stride for dimensionality reduction, the position of the only stride was shifted among the different blocks to determine its optimal setting. In particular, the best accuracies were obtained using the stride in the third residual block, as specified in Table I.

Lastly, increasing the original channel count of the stem from 64 to 96 led to enhanced accuracies, with no significant improvements beyond this size, nor using less than 64 channels.

E. Modern training procedure

An effective training procedure is as necessary as a robust architecture for achieving optimal performance. As highlighted in [50], significant advancements have been made in training methodologies for state-of-the-art neural architectures since the introduction of the **ResNet** family in 2016. Therefore, updating the **ResNeTS** training procedure to extract its full potential was crucial. Table II compares the original training procedure from **ResNet** against the improved procedure used for **ResNeTS**, which is detailed below.

The maximum number of training epochs was increased to 1500 to match InceptionTime, and an additional 10% of linear warm-up epochs were introduced at the beginning. Checkpointing was used to retain the optimal architecture upon training completion, with validation error measured every 10 epochs to guide the process. To preserve computational resources, early stopping was also triggered if no improvements in validation error were observed for a period exceeding 10% of the maximum epochs.

Inspired by the training recipe presented in [50], the Stochastic Gradient Descent (SGD) optimizer was replaced with AdamW [72], and the plateau-based learning rate scheduler was substituted with a cosine learning rate scheduler [73]. To optimize the training procedure of **ResNeTS** for the target dataset, a comprehensive hyperparameter optimization of AdamW via grid search was conducted. This resulted in a systematic exploration of all combinations of adjustable hyperparameters across extensive ranges² to ensure a thorough examination of the optimization space:

- learning rate = $\{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1\}$.
- $\beta_1/\beta_2 = \{0.5, 0.6, 0.7, 0.8, 0.9, 0.99, 0.999\}$.
- $\epsilon/\text{weight decay} = \{1 \times 10^{-8}, 1 \times 10^{-7}, 1 \times 10^{-6}, 1 \times 10^{-5}, 1 \times 10^{-4}, 1 \times 10^{-3}, 1 \times 10^{-2}, 1 \times 10^{-1}\}$.

Each combination of hyperparameters was evaluated through ten training runs of the model using a fixed set of seeds to control for randomness. This enabled the identification of the configuration that produced the best average prediction metrics

²Note that the values for β_1 , β_2 , ϵ , and weight decay were configured independently of each other.




Figure 7. Diagram of the ResNeTS architecture. A different color is assigned to each residual block.

across all trials, resulting in the values learning rate = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1 \times 10^{-6}$, weight decay = 1×10^{-3} .

Table II
COMPARISON BETWEEN THE ORIGINAL TRAINING PROCEDURE (INHERITED FROM RESNET) AND THE IMPROVED ONE IN RESNETS.

Training configuration	Original	Improved
Checkpointing	Yes	Yes
Early stopping	No	Yes
Learning rate scheduler	Reduce on plateau	Cosine
Max. training epochs	120	1500
Optimizer	SGD	AdamW
Warm-up epochs	None	150
Warm-up schedule	None	Linear

F. Ensemble of individual models

Finally, both InceptionTime and ResNeTS were run in this study as single architectures and as ensembles of 5 identical architectures, InceptionTime-5 and ResNeTS-5 respectively, for a comprehensive comparison. For ensembles, each individual architecture was initialized with different random values. Once the ensemble was trained, its prediction consisted of the mean of the individual outputs, effectively reducing the variability in the quality of predictions caused by the stochastic nature of network training.


IV. EXPERIMENTS

A. Experimental setup

1) *Competing deep learning architectures:* The capabilities of ResNeTS were compared to already established network architectures for time series analysis of remote sensing data. The first model is the MLP presented in [24] for studying the dataset used in this work. Next, the Bi-LSTM described in [34] was included to represent the capabilities of advanced RNNs. Then, the Fully Convolutional Network (FCN) described in [45] was selected to represent the early shallow CNNs used for time series analysis. The residual CNN from [45] was also included to provide a baseline of the first residual networks used for this field. Naturally, InceptionTime [53] was included to represent the state-of-the-art in residual CNNs for time series analysis. With the recent growing interest in Transformers, the architecture from [34] was included as well in the experiments. Lastly, the Rocket model [74] was employed to demonstrate the capabilities of non-deep learning models, given its high accuracy in extensive comparisons such as [54].

The MLP, Bi-LSTM, and Transformer architectures were implemented as described in the corresponding literature, since they were specifically designed for Sentinel-2 time series. Similarly, the FCN and residual CNN were used without modification due to their demonstrated adaptability across time series of various characteristics. For Rocket, although the initial implementation adhered to the default parameters for its reported adaptability, a suboptimal performance prompted an increase in the kernel count to 15000, based on its authors' advice that a higher kernel count can improve accuracy.

Original skip-connections



Upgraded skip-connections




Figure 8. Comparison between the original skip-connections and the upgraded skip-connections in [ResNeTS](#). The numbers inside each layer represent the filter size and the stride size.

Finally, adjustments were made to InceptionTime to adapt its configuration to datasets with small time series, following the guidelines of its authors [53]:

- *Bottleneck size*: The original reduction factor of $\times 0.25$ was kept to maintain the channel count throughout the network, thereby reducing the number of adjustable parameters and substantially improving runtime performance without affecting accuracy.
- *Depth*: The number of Inception modules was also kept at the default of 6, as InceptionTime authors did not observe accuracy gains when increasing the depth for short time series.
- *Filter length*: Filter sizes of $\{2, 4, 8\}$ were used, following the authors' recommendation for short time series.
- *Number of filters*: Contrary to InceptionTime authors' architectural study, an increase in the filter count to 64 in all modules resulted in precision gains in these experiments, with no further improvements beyond this point.
- *Training procedure*: To ensure a fair and direct comparison between the design principles of each architecture, the modern training procedure of [ResNeTS](#) was adopted, as it incorporates more recent techniques than InceptionTime's original configuration. A comprehensive hyperparameter optimization for AdamW was also performed on the training points, testing the same ranges of values as those used for [ResNeTS](#), resulting in selecting learning rate = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1 \times 10^{-2}$, weight decay = 1×10^{-3} .

All architectures were given a maximum of 1500 training epochs, a threshold set by InceptionTime, which had the highest epoch count among the models. In any case, checkpointing and early stopping were enabled, thereby preventing issues like overfitting to ensure optimal results for all architectures. This is particularly necessary when running numerous training epochs on smaller architectures. Within each epoch, batches of 32 elements were supplied. While the original learning rates of most architectures were proportionally adjusted to this batch size, InceptionTime and [ResNeTS](#) had their optimizers

specifically fine-tuned for these experiments, as previously described.

2) *Experimental procedure*: Each model tested, including [ResNeTS](#), was trained on the dataset described in Section II, with a single variable (Shannon, Simpson, or species richness) targeted for prediction at each instance. In particular, each combination of architecture and target variable underwent a group k-fold validation.

This k-fold process comprised five folds, each using 80% of the samples for training and 20% for testing. Additionally, 20% of the training samples were set aside in a validation set to guide periodic actions such as checkpointing, early stopping, or learning rate scheduling. Thus, the sample distribution for training, validation, and testing sets stood at 64%, 16%, and 20%, respectively. All samples were normalized so the train set followed a zero mean and a standard deviation of one.

Grouping, applied during all stages of sample division, guaranteed that observations from the same plot made in consecutive years were consistently placed in the same set of training, validation, or testing. This approach effectively mitigated temporal dependencies, in pursuit of a robust analysis. Spatial autocorrelation between observations, within and across sites, was already discarded in [24]. Species compositions and numbers strongly depend on management intensity [59], which ultimately depends on the decision of each land owner rather than on location.

For each individual fold, accuracy metrics were computed alongside architecture training time. Upon completing the k-fold process, the mean and standard deviation of these metrics were reported. Given the inherently stochastic nature of training neural architectures, the k-fold process was repeated ten times for each combination of architecture and target variable, thus recording the averaged mean and standard deviation of the metrics. To ensure a fair comparison, the same ten seeds were used in all k-folds processes, guaranteeing an identical sample distribution across architectures.

In particular, the predictive capability of each architecture was determined using the following measures: the Coefficient of Determination (r^2), the Root Mean Squared Error (RMSE), and

the Relative Root Mean Squared Error (**rRMSE**), all calculated against the mean values of the testing set. Additionally, the model-driven Systematic Root Mean Squared Error (**sRMSE**) and the data-driven Unsystematic Root Mean Squared Error (**uRMSE**) of the **RMSE** were computed to understand better the nature of the observed errors in predictions [75].

Furthermore, an explainability analysis was conducted to assess the importance of different spectral bands and time steps for the biodiversity predictions. This analysis was run on the predictions of **ResNeTS-5**, as it resulted in the most accurate architecture in this study.

The importance of spectral bands was evaluated using Shapley Additive Explanations (**SHAP**) values, derived from the DeepLift algorithm [76]. For time steps, Gradient-weighted Class Activation Mapping (**Grad-CAM**) attributions [77] were derived from the last convolutional block before the average pooling, and scaled from 8 to 16 steps. This block was selected following **Grad-CAM**'s preference of complex and abstracted features of the input data. **Grad-CAM** was used for the time steps due to the significant role it places on the order of predictors, a critical factor for temporal analysis but less significant for the bands. By combining two explainability methods, the analysis aimed to leverage their strengths while also mitigating their individual limitations, thereby providing a more robust explanation of model predictions.

Specifically, during each step of the k-fold process of **ResNeTS-5** for species richness, predictions for test data were compared to the corresponding in situ values to extract the explainability measures. For spectral bands, **SHAP** values were averaged over the temporal dimension; similarly, **Grad-CAM** attributions were averaged over the spectral dimension for time steps. These results were then averaged across the five folds of the k-fold process, and subsequently across all ten repetitions of the k-fold process, ensuring the robustness and reliability of the findings.

Additionally, to explore the internal mechanisms that allow **ResNeTS** to capture key features in the time series data, the activations of the first convolutional block were examined across all species richness experiments. To this end, within each step of the k-fold repetitions, activation maps were extracted from each convolutional filter by averaging the outputs of the convolutional block over test data. These maps highlighted which temporal features were deemed most relevant in each learned convolution across different **ResNeTS** trainings, providing insights that were validated against the **Grad-CAM** attributions.

Regarding the activation visualization, it is straightforward to do so in the first convolutional block due to its small receptive field, which allows for a more precise correlation between the activation maps and the actual time steps in the input series. In contrast, deeper convolutional blocks entangle the impact of individual time steps due to the progressive increase of the receptive field from successive convolutions and pooling operations. **Grad-CAM** internally traces gradients in the last convolutional block back to their origin, taking care of this issue.

3) *Execution environment*: All of the experiments were conducted on a computer equipped with an Intel Core i7-

11700K CPU, 128 GB of RAM, and an NVIDIA RTX 3080 Ti GPU with 12 GB of VRAM. The system ran Ubuntu 20.04, and the source code was developed on Python 3.8.0 using PyTorch 1.13.0 [78] for neural architectures, and sktime 0.15.1 [79] for Rocket. Captum 0.6.0 provided the utilities for the explainability analysis. CUDA 11.8.0 and cuDNN 8.7.0.80 [80] were leveraged to speed up execution by using single precision arithmetic. The tool GuildAI was also used for automating and tracking all experiments.

B. Results

1) *Comparison of all architectures*: Table III presents the accuracy comparison of **ResNeTS** and all the other architectures using r^2 , **rRMSE**, **sRMSE**, and **uRMSE** metrics for predicting the Shannon and Simpson indices, and species richness. Moreover, visualizations of species richness, Shannon, and Simpson predictions vs. in situ values are shown in the scatter plots of Figures 9 and 10 along with the distribution of their values.

MLP significantly underperforms all other architectures due to its inability to capture temporal and dimensional patterns between data points in time series. Note how the **RNN**, **CNNs**, and the Transformer can model these temporal and dimensional properties, thereby achieving consistently higher r^2 coefficients and lower **rRMSE** values than the **MLP**, as Table III shows.

Even though the **Bi-LSTM** demonstrates an ability to understand temporal patterns, it still falls short of the **MLP** for the Simpson index. Moreover, when compared to the **CNNs**, the **Bi-LSTM** underperforms all architectures in every experiment. As the lower accuracy is present even when compared to the shallow **FCN**, the results suggest the improved capabilities of 1D convolutions to extract more meaningful features from the time series data used.

On the one hand, the Transformer architecture improves the accuracy of the **Bi-LSTM** for both Shannon and Simpson indices and achieves similar accuracy in predicting species richness. However, **CNNs** still surpass the Transformer in accuracy, indicating that convolutions are also better suited than self-attention mechanisms for analyzing the data in this study.

Focusing on the results of the **CNNs**, it can be seen that the progressive prediction improvements by descending on the table result from a reduction in **uRMSE**, while **sRMSE** remains at a closer level to the **MLP**. The range of uncertainty of the predictions across all values is also significantly lower with the **CNNs** with respect to the **MLP**, as Figure 9 clearly shows.

Among the **CNNs**, the shallow **FCN** produces the least accurate predictions. As expected, the increased number of convolutional layers in the residual **CNNs** are more effective in extracting valuable information from the time series, thus leading to better predictions.

A comparison between all residual **CNNs** reveals **ResNeTS**'s top-accuracy, improving the results of both the residual **CNN** and InceptionTime. In particular, **ResNeTS**'s r^2 consistently achieves moderate improvements over InceptionTime's r^2 for all target variables, reaching up to +0.021 r^2 .

When constructing ensembles comprising five copies of both **ResNeTS** and InceptionTime, an enhancement in the overall




Figure 9. Scatter plots comparing the species richness predictions of the tested architectures vs. the in situ values. The predictions of each architecture are accumulated across all its experiments. (a) MLP. (b) Bi-LSTM. (c) Transformer. (d) FCN. (e) Residual CNN. (f) InceptionTime-5. (g) ResNeTS. (h) ResNeTS-5. (i) Rocket.

Table III

RECORDED ACCURACIES FOR ALL THE PREDICTED VARIABLES USING ALL THE TESTED ARCHITECTURES. INCEPTIONTIME-5 AND RESNETS-5 ARE THE ENSEMBLES OF FIVE INDIVIDUAL ARCHITECTURES. HIGHER r^2 AND LOWER RMSE INDICATE BETTER PERFORMANCES. THE BEST VALUES OF EACH COMBINATION OF VARIABLE AND METRIC ARE IN **BOLD**.

Target variable	Architecture	r^2	rRMSE	sRMSE	uRMSE
Shannon index	MLP	0.159 ± 0.042	0.234 ± 0.014	0.320 ± 0.038	0.377 ± 0.037
	Bi-LSTM	0.169 ± 0.050	0.225 ± 0.015	0.334 ± 0.057	0.335 ± 0.049
	Transformer	0.173 ± 0.063	0.215 ± 0.017	0.358 ± 0.050	0.263 ± 0.069
	FCN	0.218 ± 0.061	0.212 ± 0.013	0.316 ± 0.033	0.318 ± 0.028
	Residual CNN	0.252 ± 0.053	0.200 ± 0.013	0.318 ± 0.033	0.278 ± 0.025
	InceptionTime	0.282 ± 0.060	0.194 ± 0.013	0.321 ± 0.030	0.257 ± 0.024
	InceptionTime-5	0.300 ± 0.066	0.191 ± 0.011	0.323 ± 0.028	0.241 ± 0.023
	ResNeTS	0.283 ± 0.054	0.193 ± 0.013	0.320 ± 0.037	0.253 ± 0.019
	ResNeTS-5	0.308 ± 0.054	0.188 ± 0.012	0.321 ± 0.035	0.235 ± 0.017
Simpson index	Rocket	0.232 ± 0.086	0.201 ± 0.010	0.337 ± 0.027	0.261 ± 0.026
	MLP	0.094 ± 0.043	0.134 ± 0.013	0.086 ± 0.013	0.067 ± 0.007
	Bi-LSTM	0.088 ± 0.041	0.143 ± 0.013	0.085 ± 0.017	0.077 ± 0.012
	Transformer	0.108 ± 0.051	0.129 ± 0.017	0.094 ± 0.015	0.043 ± 0.013
	FCN	0.111 ± 0.053	0.133 ± 0.013	0.084 ± 0.012	0.068 ± 0.006
	Residual CNN	0.129 ± 0.062	0.128 ± 0.013	0.084 ± 0.012	0.061 ± 0.007
	InceptionTime	0.121 ± 0.066	0.130 ± 0.013	0.085 ± 0.012	0.062 ± 0.007
	InceptionTime-5	0.146 ± 0.066	0.126 ± 0.013	0.084 ± 0.011	0.057 ± 0.006
	ResNeTS	0.141 ± 0.058	0.125 ± 0.014	0.086 ± 0.012	0.054 ± 0.004
Species richness	ResNeTS-5	0.166 ± 0.059	0.121 ± 0.013	0.085 ± 0.012	0.049 ± 0.004
	Rocket	0.133 ± 0.065	0.126 ± 0.010	0.084 ± 0.010	0.057 ± 0.004
	MLP	0.414 ± 0.052	0.284 ± 0.024	4.469 ± 1.247	7.418 ± 0.614
	Bi-LSTM	0.512 ± 0.045	0.246 ± 0.025	4.903 ± 1.065	5.694 ± 0.545
	Transformer	0.499 ± 0.069	0.250 ± 0.025	5.504 ± 1.058	5.249 ± 0.563
	FCN	0.539 ± 0.065	0.244 ± 0.027	4.408 ± 1.103	6.028 ± 0.558
	Residual CNN	0.568 ± 0.044	0.233 ± 0.021	4.201 ± 1.034	5.753 ± 0.486
	InceptionTime	0.575 ± 0.044	0.230 ± 0.020	4.585 ± 0.920	5.352 ± 0.505
	InceptionTime-5	0.609 ± 0.035	0.222 ± 0.019	4.657 ± 0.758	4.956 ± 0.394
Species richness	ResNeTS	0.596 ± 0.039	0.223 ± 0.018	4.318 ± 0.790	5.304 ± 0.481
	ResNeTS-5	0.628 ± 0.036	0.214 ± 0.015	4.283 ± 0.790	4.968 ± 0.447
	Rocket	0.312 ± 0.047	0.289 ± 0.031	6.671 ± 1.017	5.861 ± 0.361




Figure 10. Scatter plots comparing the Shannon and Simpson indices predictions of ResNeTS-5 vs. the in situ values. The predictions of each variable are accumulated across all its experiments. The orange dashed box in the Simpson scatter plot marks a region with Simpson values above 0.5, where most of the recorded in situ values are found, and where an adequate relationship between the predicted and the actual values can be derived. (a) Shannon index. (b) Simpson index. (c) Simpson index (zoomed-in).

predictive performance is observed in both cases. Nonetheless, the **ResNeTS** ensemble continues to demonstrate competitive results compared to the InceptionTime ensemble, providing more accurate predictions for all variables.

Lastly, the Rocket regressor achieves accuracies that are lower than all the residual **CNNs**, and even worse than the **MLP** for species richness. These results contrast with studies such as [54], where Rocket can sometimes outperform advanced networks like InceptionTime. A closer look at the scatterplot in Figure 9i reveals how Rocket's predictions form a relatively wide-spread cloud of points, unlike the residual **CNNs**.

Up to this point, **ResNeTS** demonstrates a slight lead over InceptionTime in terms of accuracy. The computational performance of both architectures can become another deciding factor to further differentiate them. Table IV presents the average training time and throughput recorded for these architectures, including ensembles, aggregated across all experiments.

Table IV
COMPUTATIONAL COST OF TRAINING THE RESIDUAL **CNNs**. HIGHER SPEEDUPS ARE BETTER.

Metric	Architecture	Mean value	Speedup
Throughput (epochs/s)	InceptionTime	9.18	$\times 1.35$
	ResNeTS	12.39	
	InceptionTime-5	2.66	$\times 1.61$
	ResNeTS-5	4.29	
Train time (s)	InceptionTime	58.58	$\times 1.54$
	ResNeTS	37.94	
	InceptionTime-5	178.51	$\times 1.69$
	ResNeTS-5	105.78	

As indicated by the throughput, the lightweight architecture of **ResNeTS** achieves significantly better computational efficiency compared to InceptionTime. In particular, the individual **ResNeTS** can process 35% more samples than the individual InceptionTime within the same time frame, and the **ResNeTS-5** ensemble further increases this speedup to $\times 1.61$ over the InceptionTime-5 ensemble. A faster throughput contributes to both faster training and classification of new samples after training the architecture.

Interestingly, the training speedup of **ResNeTS** further increases to a maximum of $\times 1.69$. This implies that, in addition to the faster throughput, **ResNeTS** architectures converge in earlier epochs while training, thus triggering the early stopping sooner.

As shown, the experiments conducted demonstrate that **ResNeTS** provides state-of-the-art accuracies for a residual **CNN** in time series analysis of remote sensing data, while being significantly faster than InceptionTime in both the training and classification stages, owing to its simpler architecture.

2) *Architectural analysis of ResNeTS*: Table V displays the results of an ablation study conducted to provide empirical support for the design process outlined in Section III for developing **ResNeTS**. Specifically, starting from the final architecture, various modifications are applied to evaluate their impact on prediction quality. The training hyperparameters, such as learning rate, are optimized with each modification for a fair comparison. For brevity, the results are only reported

for species richness predictions, although similar trends were observed for Shannon and Simpson indices.

Table V
RECORDED ACCURACIES FOR THE SPECIES RICHNESS VARIABLE WHEN APPLYING DIFFERENT MODIFICATIONS TO **ResNeTS**. HIGHER r^2 VALUES ARE BETTER, WHILE LOWER **rRMSE** VALUES ARE BETTER. THE BEST VALUE OF EACH METRIC IS IN **BOLD**.

Target	Modification	r^2	rRMSE
–	Final configuration	0.596 $\pm$ 0.039	0.223 $\pm$ 0.018
Block Channels	2 nd best: 64-64-128-256	0.591 $\pm$ 0.046	0.225 $\pm$ 0.016
	3 rd best: 64-128-192-192	0.590 $\pm$ 0.040	0.225 $\pm$ 0.018
	4 th best: 64-64-128-128	0.590 $\pm$ 0.044	0.226 $\pm$ 0.019
	5 th best: 64-64-256-256	0.589 $\pm$ 0.042	0.226 $\pm$ 0.019
Block Repetitions	Two repetitions	0.568 $\pm$ 0.049	0.230 $\pm$ 0.018
Kernel Size	3	0.564 $\pm$ 0.040	0.232 $\pm$ 0.016
	7	0.591 $\pm$ 0.046	0.224 $\pm$ 0.015
	9	0.570 $\pm$ 0.048	0.230 $\pm$ 0.017
	11	0.575 $\pm$ 0.053	0.227 $\pm$ 0.016
	13	0.565 $\pm$ 0.059	0.230 $\pm$ 0.014
	15	0.583 $\pm$ 0.045	0.225 $\pm$ 0.016
Shortcut Pooling	Not included	0.582 $\pm$ 0.040	0.227 $\pm$ 0.015
Stem Channels	32	0.584 $\pm$ 0.042	0.226 $\pm$ 0.016
	64	0.585 $\pm$ 0.046	0.226 $\pm$ 0.016
	128	0.591 $\pm$ 0.038	0.224 $\pm$ 0.018
	160	0.592 $\pm$ 0.034	0.223 $\pm$ 0.014
Stride Position	Block 1	0.593 $\pm$ 0.041	0.225 $\pm$ 0.017
	Block 2	0.586 $\pm$ 0.043	0.226 $\pm$ 0.015
	Block 4	0.577 $\pm$ 0.040	0.228 $\pm$ 0.015
	No stride	0.582 $\pm$ 0.047	0.228 $\pm$ 0.017
Training Procedure	Original	0.536 $\pm$ 0.056	0.240 $\pm$ 0.022

The architectural analysis demonstrates that, regardless of the modification made, the performance of **ResNeTS** consistently declines. It's noteworthy that while some modifications lead to slight decreases in accuracy, such as the removal of pooling from the shortcuts, others result in substantial downgrades, such as retaining the original training procedure.

Overall, this study demonstrates the effectiveness of the procedure outlined in Section III in successfully accommodating **ResNet** from image analysis to time series analysis, yielding **ResNeTS**, a highly competitive architecture for the dataset of interest in this study.

3) *Importance of bands and time steps*: Figures 11 to 13 summarize the results of an explainability analysis conducted on **ResNeTS-5** to determine the importance of bands and time steps in predicting species richness. The focus on this variable offers a more robust understanding of predictor importance, given the higher accuracies of all architectures compared to Shannon and Simpson indices.

In particular, Figure 11 presents the **SHAP** values assigned to each band, while Figure 12 reflects the importance of each time step as approximated by **Grad-CAM**. The predictions of **ResNeTS** draw significantly from all features, although the architecture appears to leverage some more than others. Notably, the bands *red edge 2*, *near-infrared*, *near-infrared b*, and *short-wave infrared 2* were found to be of greatest importance. Moreover, observations from the start and the end of the growing season (April and October, respectively) proved to be highly relevant. The consistency observed in the importance across all data points underscores the reliability of the results.




Figure 11. Importance of each band in ResNeTS-5 predictions for species richness, given by SHAP values. The importance is measured and averaged across all experiments. Higher values are better.




Figure 12. Importance of each time step in ResNeTS-5 predictions for species richness, given by Grad-CAM. The importance is measured and averaged across all experiments. Higher values are better.

Although not displayed here due to space limitations, the activation maps visualized across different ResNeTS trainings on species richness align with the insights of the Grad-CAM analysis. Generally speaking, the maps show that the learned convolutional filters fall into three distinct yet non-exclusive categories: those emphasizing the beginning of the series, those focusing on the end, and those highlighting the middle portions, with the latter being significantly less common. Representative examples of different types of activation maps are displayed in Figure 13.

V. DISCUSSION

A. The motivation for ResNeTS

Although there is a wide array of literature predicting biodiversity with remote sensing data and machine learning,

recent researchers have raised awareness of their limitations due to spatial autocorrelation [19], poor domain adaptation [81], or low performance of metrics based on the spectral diversity hypothesis [22]. An approach to partially overcome these problems is to use radiative transfer models to predict the functional diversity rather than taxonomic diversity [17]. However, functional and species diversity do not always go hand in hand. Grassland communities with a high number of species may have a low functional diversity due to high functional redundancy. Thus, species richness and distribution measures such as the ones presented in this paper represent a valuable set of information for management and conservation measures.

Only a few studies make simultaneous use of the temporal and spectral dimensions to quantify taxonomic diversity; e.g., [82], [24], [14]. Other studies use only a few points in time or aggregate the temporal dimension with statistical metrics [83], [84], or aggregate the range of predicted values into a few classes [85], losing important information. Recent works still report high prediction accuracies using traditional machine learning methods without discarding spatial autocorrelation [86], which renders the use of remote sensing data potentially unnecessary [19]. [87] also pointed out the challenges encountered when modeling phenological patterns due to species diversity and image availability. Despite the success of deep learning in the field of image processing, these architectures are rarely used with full spectral-temporal data for the task of biodiversity prediction. Some architectures have shown good results when specifically fine-tuned to specific uses, such as MLPs in [24]. Although MLPs may be sufficient when predicting relatively simple biophysical variables such as above-ground biomass or LAI, more complex architectures are needed to interpret high dimensional data.

As discussed in Section I, remote sensing researchers working on tasks beyond biodiversity prediction, such as land cover classification, have already adopted advanced deep learning architectures to extend the capabilities of time series analysis [33], [34], [39], [40]. Inspired by these advancements, this work develops a powerful residual CNN to significantly improve biodiversity monitoring beyond the common MLP-based approaches.

Among the current CNNs for time series analysis, InceptionTime [53] is known for its robust performance [54]. It consists of modules that apply parallel convolutions to extract relevant features from time series data, drawing inspiration from the Inception-v4 [56] network introduced in 2017. However, most improvements in CNN development are driven by image processing, and its most advanced CNNs currently favor sequential designs over parallel ones [57], [58], [50], drawing inspiration from the success of the ResNet family [44]. This trend suggests that an effective ResNet-based CNN for time series analysis can open an interesting opportunity to advance the field by facilitating the incorporation of cutting-edge techniques from the most advanced image processing CNNs.

This observation led to the development of ResNeTS, a residual CNN resulting from adapting the ResNet architecture for time series analysis. Unlike InceptionTime, which relies on parallel convolutions with filters of varying sizes to




Figure 13. Examples of different types of activation maps extracted from the first convolutional block of **ResNeTS** in predicting species richness. (a) Map of a convolutional filter emphasizing the beginning of the series. (b) Map of a convolutional filter focusing on the end of the series. (c) Map of a convolutional filter highlighting the middle portions and the end of the series.

accommodate different time series lengths, **ResNeTS** embraces a simpler design, consisting of sequential convolutions better suited for the characteristics of the time series studied (i.e., length and dimensionality). Consequently, **ResNeTS** marks a significant step towards aligning the design principles of **CNNs** across both time series and image processing.

B. Experimental results

Empirical results demonstrate how **ResNeTS** allows the analysis of Sentinel-2 time series with improved accuracy over InceptionTime. Specifically, Shannon, Simpson, and species richness biodiversity indices are predicted, reaching gains of up to +0.02 r^2 . Furthermore, **ResNeTS** accomplishes this with fewer convolutional blocks than InceptionTime, as reflected in Figure 6b. This reduction in blocks, coupled with their light design (see Figure 6a), enable **ResNeTS** to achieve speedups of up to $\times 1.69$ in terms of execution time. This impressive speed enhancement is especially beneficial for large-scale monitoring, real-time applications, and addressing more complex time series data efficiently.

The overall better performance of **ResNeTS** over InceptionTime and the other architectures can be better understood by looking at the **uRMSE**. The **sRMSE** represents the error that can be attributed to the architecture's bias, and **ResNeTS** showed only a moderate improvement (Table III). On the other hand, the **uRMSE** represents the random error component, and its magnitude and variability across folds are consistently much lower in **ResNeTS**, compared to the other architectures. This indicates that **ResNeTS** is better at dealing with outliers and might perform consistently under scenarios with a slightly different data distribution. This is important for remote sensing studies, where unmasked haze or clouds can introduce some noise in the interpolated signal. Under circumstances with extremely bad quality pixels, the architecture returns unfeasible high species numbers that, in any case, are easy to detect and exclude during map production; i.e., above 150 species per 16 m^2 , when the maximum number recorded at the Biodiversity Exploratories is 61 (see Figure 14). Note that the Sentinel-2 time series used for calibrating the architecture was screened




Figure 14. Predictions over a poorly FORCE [63] interpolated time series outside the Biodiversity Exploratories due to cloud contamination. In these areas, the architecture returns values above 150 species per 16 m^2 , which is not plausible, as the maximum species richness value registered is 61 per 16 m^2 .

for poor-quality pixels in the Biodiversity Exploratory plots, and these extreme sources of noise do not affect the training.

It is also worth noting that, in recent years, the Transformer architecture has gained significant attention across various areas due to its competitive performance against other deep learning architectures, including **CNNs**. The Transformer evaluated in this study was also developed for handling Sentinel-2 time series [34]. Surprisingly, even basic **CNN** architectures in the experimental trials outperformed the Transformer. One potential explanation, as discussed in [51], could be the small size of the dataset used, limited by challenges such as the labor-intensive nature of in-site sample collection, and the incompatibility of aggregating data from different sources (see Section II). Other authors have already observed that Transformers are data-intensive architectures, possibly struggling in low-data regimes [88].




Figure 15. Visual differences in species richness and Shannon index predictions between **ResNetTS** and InceptionTime across various grassland patches.

When compared to results from previous works, **ResNetTS** improved the predictions of species richness from the **MLP** in [24] ($r^2 = 0.42$ vs $r^2 = 0.63$), and improved the predictions of the Shannon index to acceptable levels compared to those achieved with **MLPs** in [14] ($r^2 = 0.23$ vs. $r^2 = 0.31$). The improvement in the coefficient of determination of Simpson index predictions by **ResNetTS** with respect to the **MLP** was marginal ($r^2 = 0.16$ vs $r^2 = 0.17$). However, these errors are concentrated at low values, which correspond to plots not dominated by any particular species. This poor performance at predicting low Simpson values is probably due to the scarcity of training points for values below 0.5, as shown in Figure 10. At mid and high Simpson values (few or no dominant species) the predictions are comparable to those of other similar research [82].

C. Model interpretability

Regarding the feature importance of the Sentinel-2 bands, the results in this study are in agreement with [24]. This is,

the *red edge 2*, *near-infrared*, *near-infrared b* and *short-wave infrared 2* have the highest prediction importance for the model (Figure 11). These bands are highly sensitive to vegetation, as they can penetrate plant structures and reflect key attributes such as water and chlorophyll content. Since different plant species present distinct signatures for these characteristics, the bands enable a more accurate prediction of plant biodiversity.

Concerning the feature importance across the phenological cycle, **ResNetTS** assigns high importance values to the start and end of the growing season (April and October respectively), and lower importance towards the summer (Figure 12). This is also to some extent in agreement with [24], where May is featured as very important, rather than April. This consistency across models suggests that the relative importance of each predictor is not model-dependent, and strengthens the interpretability of the model and its results. The flowering season starts in April/May, as well as the first mowing of the most intensive grasslands (and thus species-poor plots). During October, the browning season starts, and only the extensive grasslands

(species rich) have green and senescent vegetation, whereas more intensive grasslands have either fully vegetated grass or have been recently mown (Figure 4).

It is worth noting that averaging the explainability measures over temporal and spectral dimensions poses a potential limitation of the analysis conducted. This approach, while facilitating the interpretation of the results, might mask finer dynamics, such as certain bands exhibiting more influence during particular time intervals; these subtleties are not captured in the aggregated values shown in Figure 11. Additionally, the analysis was run solely on species richness, motivated by the high accuracy of all models compared to the Shannon and Simpson indexes. Should future models achieve comparable accuracy levels across these metrics, expanding the explainability analysis could help determine whether the predictors identified as most important for species richness hold the same influence for the other biodiversity metrics.

D. Applications

Land use intensity is one of the main drivers of species richness and composition, but its effects are in some cases non-linear: while some meadows under a low land-use intensity exhibit a low species richness, some pastures under intermediate land-use intensity are home to a high number of species [3], [66]. Spatially explicit biodiversity data can contribute to a better understanding of the effects of land use on biodiversity and its relationship with ecosystem functions and services. Such data can be used to estimate other ecosystem services such as carbon storage [89] and pollination abundance [90]. Biodiversity predictions can also serve as valuable input in other models such as structure equation models, mathematic models such as Marxsan [91], or additional deep learning models [92], that attempt to disentangle the complex interactions between biodiversity, climatic and anthropogenic factors, and orient the development of conservation plans and policies.

Assessing more than one biodiversity metric across a landscape can provide a comprehensive view of biodiversity status. Figure 15 illustrates the differences in predictions between ResNeTS and InceptionTime for the Shannon and species richness variables across various grassland patches. ResNeTS predicts a broader range of values compared to InceptionTime, as expected by its higher accuracies. Interestingly, although species richness and the Shannon diversity index often correlate, the figure also demonstrates that this is not always the case, as they represent two independent components of taxonomic diversity [67].

The proposed ResNeTS offers a flexible architecture that can be applied to study not only plant biodiversity, but also any variable that can be inferred from the phenological profile of vegetation, such as crop yield prediction [37], phenometrics [12], or land use intensity [93], among many others. ResNeTS can be applied directly to other remote sensing time series without modifications, while also allowing for adjustments to achieve the most optimal performance for specific scenarios. For instance, the updated training procedure and the inclusion of pooling layers in the shortcuts are likely to provide benefits in all contexts. In contrast, in cases involving longer or higher-dimensional time series, modifications such as larger kernels,

increased convolution channels, or changes to the stride policy might be advantageous. For instance, a denser time series might be necessary to completely model the phenological profiles of other vegetation species. Naturally, ResNeTS can be tailored for classification tasks in addition to regression, by simply adapting the output layer to fit the problem at hand.

It is also worth noting how ResNeTS demonstrated its ability to effectively work with a small-sized dataset, presenting only a few hundred data points. Naturally, as with any machine learning model, it is still necessary to screen all data to minimize extremely bad quality inputs, such as pixels heavily obstructed by clouds (see Figure 14). For larger datasets, it could still be worthwhile to explore larger architectures, such as increasing the number of repetitions per block, the number of convolutional channels, or even experimenting with the larger ResNet variants than ResNet-18, which served as baseline for ResNeTS.

VI. CONCLUSIONS

This work introduced ResNeTS, a residual CNN designed for the analysis of Sentinel-2 data, and demonstrated its capabilities to enhance the prediction of biodiversity metrics, specifically species richness and Shannon and Simpson indices, with regards to other state-of-the-art architectures.

In particular, ResNeTS successfully accommodates to time series analysis the well-known ResNet computer vision architecture, and incorporates enhancements such as extended convolutional kernels for improved temporal pattern recognition, state-of-the-art training methodologies, and model ensembling to reduce the variability in predictions. These innovations allow ResNeTS to maximize the use of the spectral and temporal resolution of the time series, outperforming leading networks like InceptionTime and Transformers, in both terms of accuracy and execution time.

The advancements presented by ResNeTS extend beyond technical achievements, into biodiversity research and management. By accurately predicting biodiversity metrics across extensive regions, ResNeTS facilitates the creation of combined maps that reflect plant species richness and evenness across agricultural landscapes, and at resolutions compatible with management practices. Furthermore, integrating these predictions with climate and land-use data can help disentangle the effects of climatic and human-induced factors on biodiversity patterns, offering invaluable insights for conservation strategies and ecological studies. Finally, ResNeTS' flexibility also makes it a compelling architecture to explore in applications beyond biodiversity monitoring, such as crop yield prediction, or invasive species detection.

ACKNOWLEDGEMENTS

We thank the managers of the three exploratories, Kirsten Reichel-Jung, Iris Steitz, Sandra Weithmann, Juliane Vogt, Miriam Teuscher, and all former managers, for their work in maintaining the plot and project infrastructure; Victoria Griebmeier for giving support through the central office; Andreas Ostrowski for managing the central database; and Markus Fischer, Eduard Linsenmair, Dominik Hessenmüller,

Daniel Prati, Ingo Schöning, François Buscot, Ernst-Detlef Schulze, Wolfgang W. Weisser, and the late Elisabeth Kalko, for their role in setting up the Biodiversity Exploratories project. We thank the administration of the Hainich National Park, the UNESCO Biosphere Reserve of Schwäbische Alb, and the UNESCO Biosphere Reserve Schorfheide-Chorin, as well as all landowners, for the excellent collaboration. We also thank Ralph Bolliger for providing and maintaining the species inventory datasets, Stephan Wollawer for maintaining the remote sensing database, Florian Männer for processing the biodiversity indices, and Lisa Schwarz for the photos of plant species provided.

REFERENCES

- [1] IPBES, "Global assessment report on biodiversity and ecosystem services of the intergovernmental science-policy platform on biodiversity and ecosystem services," 2019. Version Number: 1.
- [2] R. T. Guuroh, J. C. Ruppert, J. Ferner, K. Čanak, S. Schmidtlein, and A. Linstädter, "Drivers of forage provision and erosion control in west african savannas—a macroecological perspective," *Agriculture, Ecosystems & Environment*, vol. 251, pp. 257–267, 2018.
- [3] F. Van Der Plas, "Biodiversity and ecosystem functioning in naturally assembled communities," *Biological Reviews*, p. brv.12499, 2019.
- [4] S. Díaz, J. Settele, E. S. Brondizio, H. T. Ngo, J. Agard, A. Arneeth, P. Balvanera, K. A. Brauman, S. H. M. Butchart, K. M. A. Chan, L. A. Garibaldi, K. Ichii, J. Liu, S. M. Subramanian, G. F. Midgley, P. Miloslavich, Z. Molnár, D. Obura, A. Pfaff, S. Polasky, A. Purvis, J. Razzaque, B. Reyers, R. R. Chowdhury, Y.-J. Shin, I. Visseren-Hamakers, K. J. Willis, and C. N. Zayas, "Pervasive human-driven decline of life on earth points to the need for transformative change," *Science*, vol. 366, no. 6471, p. eaax3100, 2019.
- [5] J. L. McGuire, A. M. Lawing, S. Díaz, and N. C. Stenseth, "The past as a lens for biodiversity conservation on a dynamically changing planet," *Proceedings of the National Academy of Sciences*, vol. 120, no. 7, p. e2201950120, 2023.
- [6] J. Isselstein, B. Jeangros, and V. Pavlu, "Agronomic aspects of biodiversity targeted management of temperate grasslands in europe—a review," *Agronomy research*, vol. 3, no. 2, pp. 139–151, 2005.
- [7] R. L. Schils, C. Bufe, C. M. Rhymer, R. M. Francksen, V. H. Klaus, M. Abdalla, F. Milazzo, E. Lellei-Kovács, H. T. Berge, C. Bertora, A. Chodkiewicz, C. Dămățircă, I. Feigenwinter, P. Fernández-Rebollo, S. Ghiasi, S. Hejduk, M. Hiron, M. Janicka, R. Pellaton, K. E. Smith, R. Thorman, T. Vanwallegem, J. Williams, L. Zavattaro, J. Kampen, R. Derks, P. Smith, M. J. Whittingham, N. Buchmann, and J. P. N. Price, "Permanent grasslands in europe: Land use change and intensification decrease their multifunctionality," *Agriculture, Ecosystems & Environment*, vol. 330, p. 107891, 2022.
- [8] Y. Xie, Z. Sha, and M. Yu, "Remote sensing imagery in vegetation mapping: a review," *Journal of Plant Ecology*, vol. 1, pp. 9–23, 03 2008.
- [9] C. Li, J. Xue, and B. Su, "Significant remote sensing vegetation indices: A review of developments and applications," *Journal of Sensors*, vol. 2017, p. 1353691, 2017.
- [10] J. Muro, A. Strauch, S. Heinemann, S. Steinbach, F. Thonfeld, B. Waske, and B. Dieckkrüger, "Land surface temperature trends as indicator of land use changes in wetlands," *International Journal of Applied Earth Observation and Geoinformation*, vol. 70, pp. 62–71, 2018.
- [11] C. Gómez, J. C. White, and M. A. Wulder, "Optical remotely sensed time series data for land cover classification: A review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 116, pp. 55–72, 2016.
- [12] L. Zeng, B. D. Wardlow, D. Xiang, S. Hu, and D. Li, "A review of vegetation phenological metrics extraction using time-series, multispectral satellite data," *Remote Sensing of Environment*, vol. 237, p. 111511, 2020.
- [13] S. Sesnie, C. Espinosa, A. Jara-Guerrero, and M. Tapia-Armijos, "Ensemble machine learning for mapping tree species alpha-diversity using multi-source satellite data in an ecuadorian seasonally dry forest," *Remote Sensing*, vol. 15, no. 3, p. 583, 2023.
- [14] J. Muro, A. Linstädter, F. A. Männer, L.-M. Schwarz, J. Hoffmann, and O. Dubovyk, "Predicting vegetation attributes with neural networks and sentinel-1 & 2," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLIII-B3-2022, pp. 945–950, 2022.
- [15] A. K. Schweiger and E. Laliberté, "Plant beta-diversity across biomes captured by imaging spectroscopy," *Nature Communications*, vol. 13, no. 1, p. 2767, 2022.
- [16] H. Polley, C. Yang, B. Wilsey, and P. Fay, "Spectral heterogeneity predicts local-scale gamma and beta diversity of mesic grasslands," *Remote Sensing*, vol. 11, no. 4, p. 458, 2019.
- [17] X. Ma, M. D. Mahecha, M. Migliavacca, F. van der Plas, R. Benavides, S. Ratcliffe, J. Kattge, R. Richter, T. Musavi, L. Baeten, I. Barnoiaea, F. J. Bohn, O. Bouriaud, F. Bussotti, A. Coppi, T. Domisch, A. Huth, B. Jaroszewicz, J. Joswig, D. E. Pabon-Moreno, D. Papale, F. Selvi, G. V. Laurin, F. Valladares, M. Reichstein, and C. Wirth, "Inferring plant functional diversity from space: the potential of sentinel-2," *Remote Sensing of Environment*, vol. 233, p. 111368, 2019.
- [18] J. Ferner, A. Linstädter, C. Rogass, K.-H. Südekum, and S. Schmidtlein, "Towards forage resource monitoring in subtropical savanna grasslands: going multispectral or hyperspectral?," *European Journal of Remote Sensing*, vol. 54, no. 1, pp. 364–384, 2021.
- [19] P. Ploton, F. Mortier, M. Réjou-Méchain, N. Barbier, N. Picard, V. Rossi, C. Dormann, G. Cornu, G. Viennois, N. Bayol, A. Lyapustin, S. Gourlet-Fleury, and R. Péliissier, "Spatial validation reveals poor predictive performance of large-scale ecological mapping models," *Nature Communications*, vol. 11, no. 1, p. 4540, 2020.
- [20] T. Kattenborn, J. Leitloff, F. Schiefer, and S. Hinz, "Review on convolutional neural networks (cnn) in vegetation remote sensing," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 173, pp. 24–49, 2021.
- [21] H. Zhu, Y. Huang, Y. Li, F. Yu, G. Zhang, L. Fan, J. Zhou, Z. Li, and M. Yuan, "Predicting plant diversity in beach wetland downstream of xiaolangdi reservoir with UAV and satellite multispectral images," *Science of The Total Environment*, vol. 819, p. 153059, 2022.
- [22] F. E. Fassnacht, J. Müllerová, L. Conti, M. Malavasi, and S. Schmidtlein, "About the link between biodiversity and spectral variation," *Applied Vegetation Science*, vol. 25, no. 1, p. e12643, 2022.
- [23] D. E. Rumelhart and J. L. McClelland, *Learning Internal Representations by Error Propagation*, pp. 318–362. MIT Press, 1987.
- [24] J. Muro, A. Linstädter, P. Magdon, S. Wöllauer, F. A. Männer, L.-M. Schwarz, G. Ghazaryan, J. Schultz, Z. Malenovsky, and O. Dubovyk, "Predicting plant biomass and species richness in temperate grasslands across regions, time, and land management with remote sensing and deep learning," *Remote Sensing of Environment*, vol. 282, p. 113262, 2022.
- [25] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [26] Y. Cai, K. Guan, J. Peng, S. Wang, C. Seifert, B. Wardlow, and Z. Li, "A high-performance and in-season classification system of field-level crop types using time-series landsat data and a machine learning approach," *Remote Sensing of Environment*, vol. 210, pp. 35–47, 2018.
- [27] W. Xu, W. Yang, P. Chen, Y. Zhan, L. Zhang, and Y. Lan, "Cotton fiber quality estimation based on machine learning using time series uav remote sensing data," *Remote Sensing*, vol. 15, no. 3, 2023.
- [28] R. J. Williams and D. Zipser, "A Learning Algorithm for Continually Running Fully Recurrent Neural Networks," *Neural Computation*, vol. 1, pp. 270–280, 06 1989.
- [29] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, pp. 2278–2324, Nov 1998.
- [30] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, pp. 1735–1780, 11 1997.
- [31] K. Cho, B. van Merriënboer, Ç. Gülçehre, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," *CoRR*, vol. abs/1406.1078, 2014.
- [32] H. Ma and S. Liang, "Development of the glass 250-m leaf area index product (version 6) from modis data using the bidirectional lstm deep learning model," *Remote Sensing of Environment*, vol. 273, p. 112985, 2022.
- [33] H. Crisóstomo de Castro Filho, O. Abílio de Carvalho Júnior, O. L. Ferreira de Carvalho, P. Pozzobon de Bem, R. dos Santos de Moura, A. Olino de Albuquerque, C. Rosa Silva, P. H. Guimarães Ferreira, R. Fontes Guimarães, and R. A. Trancoso Gomes, "Rice crop detection using lstm, bi-lstm, and machine learning models from sentinel-1 time series," *Remote Sensing*, vol. 12, no. 16, 2020.
- [34] M. Rußwurm and M. Körner, "Self-attention for raw optical satellite time series classification," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 169, pp. 421–435, 2020.
- [35] A. Garioud, S. Valero, S. Giordano, and C. Mallet, "Recurrent-based regression of sentinel time series for continuous vegetation monitoring," *Remote Sensing of Environment*, vol. 263, p. 112419, 2021.

- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), vol. 30, Curran Associates, Inc., 2017.
- [37] J. Xu, Y. Zhu, R. Zhong, Z. Lin, J. Xu, H. Jiang, J. Huang, H. Li, and T. Lin, "Deepcropmapping: A multi-temporal deep learning approach with improved spatial generalizability for dynamic corn and soybean mapping," *Remote Sensing of Environment*, vol. 247, p. 111946, 2020.
- [38] Z. Lin, R. Zhong, X. Xiong, C. Guo, J. Xu, Y. Zhu, J. Xu, Y. Ying, K. C. Ting, J. Huang, and T. Lin, "Large-scale rice mapping using multi-task spatiotemporal deep learning and sentinel-1 sar time series," *Remote Sensing*, vol. 14, no. 3, 2022.
- [39] H. Wang, X. Zhao, X. Zhang, D. Wu, and X. Du, "Long time series land cover classification in china from 1982 to 2015 based on bi-lstm deep learning," *Remote Sensing*, vol. 11, no. 14, 2019.
- [40] Y. Xi, C. Ren, Q. Tian, Y. Ren, X. Dong, and Z. Zhang, "Exploitation of time series sentinel-2 data and different machine learning algorithms for detailed tree species classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 7589–7603, 2021.
- [41] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, p. 84–90, may 2017.
- [42] H. Zhao, Z. Chen, H. Jiang, W. Jing, L. Sun, and M. Feng, "Evaluation of three deep learning models for early crop classification using sentinel-1a imagery time series—a case study in zhanjiang, china," *Remote Sensing*, vol. 11, no. 22, 2019.
- [43] L. Zhong, L. Hu, and H. Zhou, "Deep learning based multi-temporal crop classification," *Remote Sensing of Environment*, vol. 221, pp. 430–443, 2019.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, June 2016.
- [45] Z. Wang, W. Yan, and T. Oates, "Time series classification from scratch with deep neural networks: A strong baseline," in *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 1578–1585, May 2017.
- [46] G. Paolini, M. J. Escorihuela, O. Merlin, M. P. Sans, and J. Bellvert, "Classification of different irrigation systems at field scale using time-series of remote sensing data," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 10055–10072, 2022.
- [47] T. Lin, Y. Wang, X. Liu, and X. Qiu, "A survey of transformers," *AI Open*, vol. 3, pp. 111–132, 2022.
- [48] X. He, Y. Zhou, J. Zhao, D. Zhang, R. Yao, and Y. Xue, "Swin transformer embedding unet for remote sensing image semantic segmentation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–15, 2022.
- [49] S. K. Roy, A. Deria, D. Hong, B. Rasti, A. Plaza, and J. Chanussot, "Multimodal fusion transformer for remote sensing image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–20, 2023.
- [50] Z. Liu, H. Mao, C. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A convnet for the 2020s," *CoRR*, vol. abs/2201.03545, 2022.
- [51] Y. Liu, E. Sangineto, W. Bi, N. Sebe, B. Lepri, and M. Nadai, "Efficient training of visual transformers with small datasets," *Advances in Neural Information Processing Systems*, vol. 34, pp. 23818–23830, 2021.
- [52] T. A. W. Aaron E. Maxwell and F. Fang, "Implementation of machine-learning classification in remote sensing: an applied review," *International Journal of Remote Sensing*, vol. 39, no. 9, pp. 2784–2817, 2018.
- [53] H. Ismail Fawaz, B. Lucas, G. Forestier, C. Pelletier, D. F. Schmidt, J. Weber, G. I. Webb, L. Idoumghar, P.-A. Muller, and F. Petitjean, "Inceptiontime: Finding alexnet for time series classification," *Data Mining and Knowledge Discovery*, vol. 34, no. 6, pp. 1936–1962, 2020.
- [54] A. P. Ruiz, M. Flynn, J. Large, M. Middlehurst, and A. Bagnall, "The great multivariate time series classification bake off: a review and experimental evaluation of recent algorithmic advances," *Data Mining and Knowledge Discovery*, vol. 35, no. 2, pp. 401–449, 2021.
- [55] M. Rußwurm, C. Pelletier, M. Zollner, S. Lefèvre, and M. Körner, "Breizhcrops: A time series dataset for crop type mapping," *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences ISPRS (2020)*, 2020.
- [56] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, Feb. 2017.
- [57] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, "Searching for mobilenetv3," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [58] M. Tan and Q. Le, "Efficientnetv2: Smaller models and faster training," in *Proceedings of the 38th International Conference on Machine Learning* (M. Meila and T. Zhang, eds.), vol. 139 of *Proceedings of Machine Learning Research*, pp. 10096–10106, PMLR, 18–24 Jul 2021.
- [59] N. Blüthgen, C. F. Dormann, D. Prati, V. H. Klaus, T. Kleinebecker, N. Hölzel, F. Alt, S. Boch, S. Gockel, A. Hemp, J. Müller, J. Nieschulze, S. C. Renner, I. Schöning, U. Schumacher, S. A. Socher, K. Wells, K. Birkhofer, F. Buscot, Y. Oelmann, C. Rothenwöhrer, C. Scherber, T. Tschamtker, C. N. Weiner, M. Fischer, E. K. Kalko, K. E. Linsenmair, E.-D. Schulze, and W. W. Weisser, "A quantitative index of land-use intensity in grasslands: Integrating mowing, grazing and fertilization," *Basic and Applied Ecology*, vol. 13, no. 3, pp. 207–220, 2012.
- [60] M. Fischer, O. Bossdorf, S. Gockel, F. Hänsel, A. Hemp, D. Hessenmöller, G. Korte, J. Nieschulze, S. Pfeiffer, D. Prati, S. Renner, I. Schöning, U. Schumacher, K. Wells, F. Buscot, E. K. Kalko, K. E. Linsenmair, E.-D. Schulze, and W. W. Weisser, "Implementing large-scale and long-term functional biodiversity research: The biodiversity exploratories," *Basic and Applied Ecology*, vol. 11, no. 6, pp. 473–485, 2010.
- [61] C. J. Keylock, "Simpson diversity and the shannon-wiener index as special cases of a generalized entropy," *Oikos*, vol. 109, no. 1, pp. 203–207, 2005.
- [62] R. Bolliger, D. Prati, and M. Fisher, "Vegetation records for grassland EPs, 2008 - 2020. biodiversity exploratories information system (bexis). dataset ID = 27386," 2021.
- [63] D. Frantz, "Force—landsat + sentinel-2 analysis ready data and beyond," *Remote Sensing*, vol. 11, no. 9, 2019.
- [64] E. S. Agency, "Spectral resolution – sentinel online," 2023. Accessed: September 25, 2023.
- [65] L. Ma, Y. Liu, X. Zhang, Y. Ye, G. Yin, and B. A. Johnson, "Deep learning in remote sensing applications: A meta-analysis and review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 152, pp. 166–177, 2019.
- [66] V. Busch, V. H. Klaus, D. Schäfer, D. Prati, S. Boch, J. Müller, M. Chisté, K. Mody, N. Blüthgen, M. Fischer, N. Hölzel, and T. Kleinebecker, "Will i stay or will i go? plant species-specific response and tolerance to high land-use intensity in temperate grassland ecosystems," *Journal of Vegetation Science*, vol. 30, no. 4, pp. 674–686, 2019.
- [67] A. E. Magurran, "Measuring biological diversity," *Current Biology*, vol. 31, no. 19, pp. R1174–R1177, 2021.
- [68] W. E. Kunin, J. Harte, F. He, C. Hui, R. T. Jobe, A. Ostling, C. Polce, A. Šizling, A. B. Smith, K. Smith, S. M. Smart, D. Storch, E. Tjørve, K.-I. Ugland, W. Ulrich, and V. Varma, "Upscaling biodiversity: estimating the species–area relationship from small samples," *Ecological Monographs*, vol. 88, no. 2, pp. 170–187, 2018.
- [69] W. E. Kunin, J. Harte, F. He, C. Hui, R. T. Jobe, A. Ostling, C. Polce, A. Šizling, A. B. Smith, K. Smith, S. M. Smart, D. Storch, E. Tjørve, K. Ugland, W. Ulrich, and V. Varma, "Upscaling biodiversity: estimating the species–area relationship from small samples," *Ecological Monographs*, vol. 88, no. 2, pp. 170–187, 2018.
- [70] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, June 2009.
- [71] T. He, Z. Zhang, H. Zhang, Z. Zhang, J. Xie, and M. Li, "Bag of tricks for image classification with convolutional neural networks," *CoRR*, vol. abs/1812.01187, 2018.
- [72] I. Loshchilov and F. Hutter, "Fixing weight decay regularization in adam," *CoRR*, vol. abs/1711.05101, 2017.
- [73] I. Loshchilov and F. Hutter, "SGDR: stochastic gradient descent with restarts," *CoRR*, vol. abs/1608.03983, 2016.
- [74] A. Dempster, F. Petitjean, and G. I. Webb, "Rocket: exceptionally fast and accurate time series classification using random convolutional kernels," *Data Mining and Knowledge Discovery*, vol. 34, pp. 1454–1495, Sep 2020.
- [75] C. J. Willmott, "On the validation of models," *Physical Geography*, vol. 2, no. 2, pp. 184–194, 1981.
- [76] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), vol. 30, Curran Associates, Inc., 2017.
- [77] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via

gradient-based localization,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.

- [78] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Z. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” *CoRR*, vol. abs/1912.01703, 2019.
- [79] M. Löning, A. J. Bagnall, S. Ganesh, V. Kazakov, J. Lines, and F. J. Király, “sktime: A unified interface for machine learning with time series,” *CoRR*, vol. abs/1909.07872, 2019.
- [80] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer, “cudnn: Efficient primitives for deep learning,” *CoRR*, vol. abs/1410.0759, 2014.
- [81] H. Meyer and E. Pebesma, “Machine learning-based global maps of ecological variables and the challenge of assessing them,” *Nature Communications*, vol. 13, no. 1, p. 2208, 2022.
- [82] M. Fauvel, M. Lopes, T. Dubo, J. Rivers-Moore, P.-L. Frison, N. Gross, and A. Ouin, “Prediction of plant diversity in grasslands using sentinel-1 and -2 satellite image time series,” *Remote Sensing of Environment*, vol. 237, p. 111536, 2020.
- [83] H. Gholizadeh, J. A. Gamon, C. J. Helzer, and J. Cavender-Bares, “Multi-temporal assessment of grassland α - and β -diversity using hyperspectral imaging,” *Ecological Applications*, vol. 30, no. 7, p. e02145, 2020.
- [84] S. Bae, S. R. Levick, L. Heidrich, P. Magdon, B. F. Leutner, S. Wöllauer, A. Serebryanyk, T. Nauss, P. Krzystek, M. M. Gossner, P. Schall, C. Heibl, C. Bässler, I. Doerfler, E.-D. Schulze, F.-S. Krah, H. Culmsee, K. Jung, M. Heurich, M. Fischer, S. Seibold, S. Thorn, T. Gerlach, T. Hothorn, W. W. Weisser, and J. Müller, “Radar vision in the mapping of forest biodiversity from space,” *Nature Communications*, vol. 10, no. 1, p. 4757, 2019.
- [85] R. A. Crabbe, D. Lamb, and C. Edwards, “Discrimination of species composition types of a grazed pasture landscape using sentinel-1 and sentinel-2 data,” *International Journal of Applied Earth Observation and Geoinformation*, vol. 84, p. 101978, 2020.
- [86] A. Safonova, G. Ghazaryan, S. Stiller, M. Main-Knorn, C. Nendel, and M. Ryo, “Ten deep learning techniques to address small data problems with remote sensing,” *International Journal of Applied Earth Observation and Geoinformation*, vol. 125, p. 103569, Dec. 2023.
- [87] N. Younes, K. E. Joyce, and S. W. Maier, “All models of satellite-derived phenology are wrong, but some are useful: A case study from northern australia,” *International Journal of Applied Earth Observation and Geoinformation*, vol. 97, p. 102285, 2021.
- [88] Z. Lu, H. Xie, C. Liu, and Y. Zhang, “Bridging the gap between vision transformers and convolutional neural networks on small datasets,” in *Advances in Neural Information Processing Systems* (A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, eds.), 2022.
- [89] Y. Yang, D. Tilman, G. Furey, and C. Lehman, “Soil carbon sequestration accelerated by restoration of grassland biodiversity,” *Nature Communications*, vol. 10, p. 718, Feb. 2019.
- [90] H. Feilhauer, D. Doktor, S. Schmidtlein, and A. K. Skidmore, “Mapping pollination types with remote sensing,” *Journal of Vegetation Science*, vol. 27, no. 5, pp. 999–1011, 2016.
- [91] I. Ball, H. Possingham, and M. Watts, *Spatial conservation prioritisation: Quantitative methods and computational tools. Chapter 14: Marxan and relatives: Software for spatial conservation prioritisation*. Oxford University Press, 2009.
- [92] D. Silvestro, S. Gorla, T. Sterner, and A. Antonelli, “Improving biodiversity protection through artificial intelligence,” *Nature Sustainability*, vol. 5, no. 5, pp. 415–424, 2022.
- [93] M. Lange, H. Feilhauer, I. Kühn, and D. Doktor, “Mapping land-use intensity of grasslands in germany with machine learning and sentinel-2 time series,” *Remote Sensing of Environment*, vol. 277, p. 112888, 2022.



Álvaro G. Dieste received a B.Sc. in Computer Engineering and an M.Sc. in High Performance Computing in 2021 and 2022, respectively, from the University of Santiago de Compostela, Santiago, Spain, where he is currently working towards a Ph.D. in Computer Science.

He is currently a Predoctoral Researcher at the Singular Research Center on Intelligent Technologies (CiTIUS) at the University of Santiago de Compostela. His research interests focus on developing computer vision techniques for vegetation monitoring using remote sensing data. He integrates knowledge from deep learning, high performance computing, and big data to advance his field.



Francisco Argüello received a B.Sc. in Physics and a Ph.D. in Physics in 1988 and 1992, respectively, from the University of Santiago de Compostela, Santiago, Spain.

He is currently a Full Professor in the Department of Electronic and Computer Engineering at the University of Santiago de Compostela. His research interests include signal and image processing, computer graphics, parallel and distributed computing, and quantum computing.



Dora B. Heras received an M.Sc. in Physics and a Ph.D. in Physics in 2000, from the University of Santiago de Compostela, Santiago, Spain. Since 2023 she has been co-chair of the International Euro-Par conference. Since 2020, she has held the position of Chair for the High-Performance and Disruptive Computing in Remote Sensing (HDCRS) Working Group under the IEEE GRSS Earth Science Informatics Technical Committee (ESI TC).

She is currently a Full Professor in the Department of Electronics and Computer Engineering at the University of Santiago de Compostela. Her research interests cover a range of topics in the combined fields of image processing, remote sensing, machine learning, and high performance computing. In particular, she has published papers on high performance computing, registration, classification, domain adaptation, and change detection applied to remotely sensed images.



Paul Magdon received a B.Sc. in 2005, an M.Sc. in 2008, and a Ph.D. in Forest Science in 2013 from the University of Göttingen, Germany.

In 2022, he started as a professor for geoinformation and forest planning at the University of Applied Sciences and Arts (HAWK) in Göttingen. His research focuses on the development and optimization of remote sensing and GIS applications for monitoring natural resources. He is working on developing data- and sensor-driven environmental monitoring systems at the plot, regional, and national levels. The combination of field observations, geospatial analysis, and data from various sensors into coherent monitoring networks is one of his specific interests.



Olena Dubovyk received a Ph.D. from the University of Bonn in 2014. Additionally, she possesses a joint M.Sc. in Geo-information Science and Earth Observation for Environmental Modelling and Management from the University of Southampton in Great Britain, the Lund University in Sweden, and the University of Twente in the Netherlands. She also has an M.Sc. and a B.Sc. in Geography from the Taras Shevchenko National University of Kyiv, Ukraine.

She is currently a Professor at the University of Hamburg. She is involved in the field of remote sensing, particularly in environmental monitoring and management. Her expertise extends to using Earth observation and geo-information systems for various applications including drought monitoring and assessment, as well as supporting policy and international cooperation (e.g., SDGs, Sendai indicators) and disaster risk reduction.



Anja Linstädter received an M.Sc. in Biology in 1994 from the University of Hamburg, and a Ph.D. in 2001 from the University of Cologne.

She is currently a full Professor of Biodiversity Research and Systematic Botany at the Institute of Biochemistry and Biology, University of Potsdam, Germany. As a range ecologist, she investigates the impacts and adaptations of grasslands under global changes. She has led several German and European-funded projects across African and European landscapes.



Javier Muro received a B.Sc. in 2009 from the University of Alcalá, Spain, an M.Sc. in 2014 from the University of Turku, Finland, and a Ph.D. in remote sensing in 2019 from the University of Bonn, Germany.

Since 2022 he has worked as a research associate at the Institute of Farm Economics of the Thünen Institute, Germany. His expertise focuses on the use of Earth Observation, machine learning, and data science to study the interactions between biodiversity, ecosystem services, and land management. He has

led multiple publications on modeling different land use processes in Latin America, Europe, and East Africa.