

Detecting the necessity of forest thinning with deep learning: An application in Lungau region, Austria

Philipp Satlawa and Robert B. Fisher

Abstract

Timely information about the necessity of thinning in the forest is vital for forest management to maintain a healthy forest while maximizing income. Currently, very high spatial resolution remote sensing data can provide crucial assistance for the experts to evaluate the maturity of thinnings. Yet, this task is still predominantly determined in the field and demands extensive resources. This paper presents a deep convolutional neural network (DCNN) to detect the necessity and urgency of thinnings by using only remote sensing data. The approach uses very high spatial resolution RGB and near-infrared orthophotos, canopy height model (CHM), digital terrain model (DTM), slope and the reference data, in this case from spruce dominated forests in the Austrian Alps. After tuning, the model achieves a test set F1 score of 82.23%, which is practically usable. We conclude that DCNNs are capable of detecting the need for thinning in forests. In contrast, attempts of assessing the urgency of thinnings with DCNNs proved to be unsuccessful. However, additional data such as age or yield class has the potential of improving the results. Investigation of the influence of the individual input features showed that orthophotos appear to contain the most relevant information for detecting the demand for thinning. Moreover, we observe a gain in performance by adding CHM and slope, whereas adding the DTM harms the model's performance.

Keywords: forest management, agroforestry, computer vision, deep learning, thinning

1 Introduction

Maintaining a healthy, stable forest to produce valuable wood requires fostering the forest. An essential technique to accomplish this is a silvicultural operation called thinning [12]. The main objective of thinning is to regulate the vertical space of trees, thus steering the allocation of the available resources (e.g. sunlight, water, nutrients) into the stems of remaining higher quality trees [48]. Although the primary objective of thinning is to prepare the forest stand for the final harvest, it also provides forest owners with the opportunity of obtaining additional income by selling the removed trees [43, 58, 31, 55].

Determining if a forest stand needs thinning is a complex task since it depends on many factors such as soil quality, age and tree species composition [37]. This assessment makes assessing a forest stand for the necessity of thinning a non-trivial task, and this is why the job is still mainly conducted by specialized forest personnel. Sending personnel into the field is often very expensive. Hence, thinning assessment is either not done at all, like in the case of small forest owners or done at long time intervals. Long time intervals might be sufficient for slowly growing sites, however vital areas with high mean annual increment are often overlooked, which results in a sub-optimal timing of thinning in these stands.

Recently, the proliferation in the amount and enhancement in the quality of remote sensing data has raised the possibility of providing uniform and detailed aerial information over large areas [17]. Furthermore, with the advances in sensor technology as well as computing power, the interest in applying algorithms for automated extraction of forest parameters from remotely sensed spectral information has increased [44]. Aerial forestry data coupled with recent increases in computing power allows machine learning (ML) algorithms to be deployed, such as for change

detection in forests [32, 66], classifying tree species [4, 14, 34, 3] or estimating the standing volume [20, 6, 15].

Despite the high importance of timely planning of forest operations, few studies have been conducted to predict the necessity of thinning from remote sensing data. Just two known studies addressed the challenge. First, [33] tried to predict forest operations at the stand-level by using Landsat-TM satellite imagery with moderate success. Then [63] employed ALS data-derived features to predict the thinning maturity at the stand-level. This study showed much better results, with classification accuracy ranging from 79% to 83% for predicting the timing of the subsequent thinning. Another approach to predict the necessity of thinning is to utilize key forest parameters acquired through remote sensing together with additional inventory data to create statistical models [19]. Nevertheless, ALS data is still expensive to obtain, and just a few countries acquire it systematically.

To extract information about the necessity of thinnings from widespread areal images, we use semantic segmentation, as it was successfully used for many remote sensing tasks [40]. Semantic segmentation is a computer vision task where the algorithm labels each pixel or patch of an image to a predefined range of classes. Accelerated by the first end-to-end fully convolutional network (FCN) [42], the utilization of end-to-end DCNNs for the task of semantic segmentation increased greatly. Since then, many different architectures were proposed and adapted to remote sensing applications [64, 67, 13]. The main advantage of using DCNNs for semantic segmentation is their effectiveness in extracting complex features from wide receptive fields. Nonetheless, this capability comes with the price of not being able to maintain high spatial resolution and results in inaccurate and blurred boundaries between the classes. To counteract this, newer DCNNs use more pronounced/distinct encoder-decoder architectures with skip connections. For example, UNet applies such skip connections in its symmetrical encoder-decoder architecture [52], where the features extracted in the encoder are directly coupled to the corresponding decoder layers. The outstanding performance results, together with the simplicity of the architecture, ensured the wide adoption in the remote sensing research community for a variety of applications such as ship detection [28], road extraction [11] and land cover classification [60].

To obtain models that can learn more complex input representations, deeper DCNNs with more stacked layers were created. However, these deeper DCNNs often resulted in worse-performing models than the more shallow predecessors due to the degradation problem [21]. This issue was resolved by Deep residual networks (e.g. ResNet) that employed residual blocks, which added an identity shortcut connection and overcame the degradation problem [24]. Further development of the DCNNs resulted in the creation of a network architecture called DenseNet, which utilizes dense blocks that introduce direct connections from any layer to all subsequent layers to improve further the information flow between layers [30]. [36] adopted the DenseNet connection structure for semantic segmentation by applying dense blocks to the UNet like symmetric encoder-decoder structure. This merge resulted in a DCNN called FC-DenseNet that needs fewer parameters while performing better on various semantic segmentation challenges.

Another compelling approach to resolve the trade-off between high context extraction with heavy downsampling and accurate boundary prediction is the use of dilated or atrous convolutions. Atrous convolutions contribute with a convolution filter that is spaced apart, thus attaining a wider field of view while retaining its spatial dimension. [8] proposed a DCNN architecture called DeepLabv1 that incorporates the atrous convolutions in a VGG-16 architecture to address the trade-off between high context extraction with heavy downsampling. After several revisions of this architecture that included, among other things, a Spatial Pyramidal Pooling that was introduced in SPPNet [23] and adopted by [9] to create the Atrous Spatial Pyramid Pooling (ASPP) in DeepLabv2. The current network architecture DeepLabv3+ [10], provides some of the best results in semantic segmentation challenges. Furthermore, its use in remote sensing applications is auspicious. For example, [41] showed the network’s excellent performance in classifying marsh vegetation in China. Accordingly, DeepLabv3+ is the network used in Section 3.2 to classify tree thinning regions.

Nevertheless, little research has been done to detect the necessity of thinning solely with remote sensing data. This paper presents a model capable of achieving this task.

In this study, we tested a deep learning approach to determine if remote sensing data alone could be used to detect the need for thinning. The two factors investigated were whether the approach could:

1. Detect if thinning is needed, based on deep convolutional neural networks (DCNNs) trained on very high spatial resolution imagery, and
2. Differentiate between thinnings with different urgencies (i.e. timescales).

An accurate prediction of thinnings from remote sensing data has the potential of delivering information for vast or remote forest areas promptly and thus improving the stability and wood quality of forests.

2 Previous Research and Background

This section provides background information related to detecting the necessity of thinnings with the help of remote sensing. First, we introduce thinning, its effect on the forest stand, and its impact on the timber quality. After that, we show the current state of remote sensing in forestry on its most important research questions. Finally, we overview the current state of the art on image change analysis techniques and its application in remote sensing.

2.1 Thinning

Thinning can be seen as the primary steering technique in preparation for the final harvest. Although many types of thinning have emerged through the history of forestry, the main objective of all of them is to remove some trees to accelerate the growth in the remaining future crop trees. Notably, this enables the remaining trees to allocate the newly available resources mainly into their basal stem [48, 26], which is economically the most valuable part of the tree.

Furthermore, by applying selective thinning, hence selecting the most promising future crop trees, log quality can be significantly enhanced [59, 45]. Thus, thinning provides an opportunity to increase the amount of good quality timber in the final harvest while producing additional income from the sale of the removed trees. Since the main objective of many forest owners is to maximize the net present value of their forest area, thinned stands outperform unmanaged stands [58, 31, 25].

One of the main concerns of performing thinning is the enhanced risk for damage to the forest due to high wind or snow. Various studies have shown that right after thinning, a forest stand's stability is lower due to the higher roughness of the tree crowns [50]. Although this is true for the time right after thinning, the risk of injuries decreases with time until no additional risk is present anymore [50, 46]. Nonetheless, thinning in Norway spruce stands has been recorded to reduce the h/d ratio of trees (height to diameter ratio) [56]. The h/d ratio is an indicator for the individual stability of a tree, where low h/d values are equivalent to high tree stability [51, 61, 62]. Thus, newer research suggests that early heavy thinning might even increase the stability of a stand [54]. Thereby, timeliness is crucial, as delaying thinnings results in a risk increase of damage in the stands [51, 7].

Optimal thinning schemes differ by the tree species present, any previously applied measures, and the density of the standing trees. Due to this variation, it is essential to define the type of thinning that is being used throughout this study. The study area is predominantly stocked with Norway spruce dominated coniferous forest. [1] and later [25] found that the highest economic return on investment is produced by selective crown thinning with a selection of target trees for this typical forest type. The Austrian Federal Forest adopted this thinning type as the standard thinning scheme and thus, nearly all thinnings are planned as such. Therefore, when we refer to thinning in this study, it is equivalent to this specific thinning type.

Thinnings of Norway Spruce stands of the Austrian Federal Forest are executed 2-3 times during a lifetime of a stand, where 1/4 to 1/3 of the actual stocking volume is being harvested

each time. The first thinning is carried out typically at around 13-19 m top height, followed by a second thinning at approximately 20-30 m top height, with a third thinning on good-growing sites. The exact timing of the procedure depends on the density of the standing trees, which is related to conditions such as the timing of the previous thinning and the growth performance at the site.

2.2 Remote sensing in forestry

Utilizing remote sensing imagery for extracting information about the forest structure and its employment in forest planning has been practised since the 1960s [2]. Remote sensing provides the opportunity to gather uniform information about larger areas otherwise impossible to collect from field measurements. Automated derivation of crucial information from the forest, such as tree species classification and wood volume estimation, has long been a goal of researchers in forestry [38]. Traditionally, collecting information about the forest structure was based on sending experts into the field as in conventional forest inventories. This process is costly and often not affordable for small forest owners. With the advances in sensor technology as well as computing power, the interest in applying algorithms for automated extraction of forest parameters from remotely sensed spectral information increased [44].

At present, the leading sensing instrument technologies applied for retrieving forestry relevant metrics are multi-spectral or hyperspectral cameras, light detection and ranging (lidar), and to a certain extent, Synthetic-aperture radar (SAR) [27]. These instruments can be mounted on three different platforms, satellite, aircraft and unmanned aerial vehicle (UAV), each giving the sensing instrument different ranges of spatial resolution. Moreover, sensors mounted on a plane can support conventional forest inventories and provide helpful information for operational forest management [47], [38].

Despite the importance of forest operations such as thinning, few studies have tried to create models that can predict the necessity of thinning directly from remote sensing data [63]. Most research has been focused on change detection, tree species classification as well as wood volume estimation. Although substantial progress has been made in all the introduced research questions, they are far from being solved and still active research subjects.

2.3 Change Analysis in Forest Management

A forest changes continuously, whether through carried out forest operations such as thinning or clear-cutting or by forest damage caused by heavy winds, fire and other natural disasters. The man-made changes can be reported and easily updated. However, the changes induced by nature must be spotted differently. Automatic detection of changes in the forest using aerial imagery shows promising results in areas where severe changes have happened, such as clear-cutting and intense storm damage. However, moderate changes such as thinning are much harder to discover [32]. When using airborne laser scanning (ALS) as the data source to detect changes, [66] were able to identify the removal of individual trees. Nonetheless, acquiring aerial images or even more expensive ALS data is often done at considerable intervals. Hence, a more interesting remote sensing platform for this problem are satellites. Due to their higher temporal resolution, satellites can provide much more timely data as it is often required in disaster response. In particular, SAR data can provide data even on overcast days. Although weather conditions influence the signal, this noise can be filtered, as [49] has demonstrated. Thereby they were able to detect changes in forest areas greater than 1 ha with Sentinel-1 data.

3 Materials and Methods

3.1 Materials

For training the machine learning model, we used data collected for the Lungau region through aerial imagery, airborne laser scanning (ALS) as well as data from forest management plans. This

data was first cleaned and preprocessed before being used as input data for the machine learning algorithms, resulting in 5 distinct input types as illustrated in Table 1. This section defines what kind of thinning is used in this study, describes the data and its acquisition, and the data preprocessing.

3.1.1 Study area

Located in the Lungau region (Tamsweg district), Austria (47°00' - 47°13'N, 13°23' - 14°0'E, UTM/WGS84 projection), the study area is managed by the Austrian Federal Forests (ÖBF AG). As part of the Central Eastern Alps, the area is mountainous with an altitude range 993 – 1906 m and has a distinct alpine climate with an average annual temperature of 5.2 °C, mean annual precipitation ranging between 770 – 840 mm and a snow cover minimum of 1 cm on 105 days per year. The study area illustrated in Figure 1 has an area of 21826.55 ha and is predominantly covered with forest (63.9% of the area). The forest area is further divided into commercial forest and protective forest. The commercial forest is managed to maximize the income from timber production while minimizing the risk of forest damage. In contrast, a protective forest objective is to protect against avalanches, rockfall, erosion and floods. Although thinnings are planned in both forest types, thinnings in the protective forest aim to ensure its protective function, whereas thinnings, as specified in section 2.1, are just planned in a commercial forest. We are merely interested in the part where commercial thinnings are feasible. Hence we restrict the study area to the commercial forest. The commercial forest has an extent of 9353.54 ha and is stocked with mainly coniferous forest. The most frequent tree species are 81.0% Norway spruce (*Picea abies*, 81.0%) and European larch (*Larix decidua*, 17.6%). The remaining 1.4% of the area are stocked with deciduous forest.

3.1.2 Remote sensing and ground truth data acquisition

As shown in Table 1, all data came from three sources and was acquired at three disparate points in time. Ideally, all data should be acquired simultaneously, but we were limited in this study.

Table 1: Data sources. NIR: near infrared, CHM: Canopy height model, DTM: digital terrain model, AI: aerial imagery, ALS: airborne laser scanning, RD: actual thinning reference data, Res.: spatial resolution, S.E.: standard error and Year: acquisition year. Slope data computed from the DTM

Name	Source	Format	Bands	Res. [m]	S.E. [m]	Year
RGB + NIR	AI	GeoTIFF	4	0.2	-	2018
CHM	AI	GeoTIFF	1	1	1	2018
DTM	ALS	GeoTIFF	1	1	0.15	2013
Slope	ALS	GeoTIFF	1	1	-	2013
Ground Truth	RD	Shape	3	-	-	2017

The acquisition of the aerial imagery is performed by the States of Austria, where one-third of Austria is updated every year. All aerial imagery for the study site was recorded on 11.09.2018, and 12.09.2018 with an UltraCam Eagle Mark 3 431S61680X916102-f100 mounted on a Beechcraft Super King Air B200 D-IWAW. Four bands were acquired with a spatial resolution of 20 cm, the three bands for RGB and one band in the near-infrared (NIR) spectrum (Table 1). After the acquisition, the raw data was geometrically and radiometrically corrected with the corresponding calibration data of the camera and the four channels were stitched together using the monolithic stitching method [18]. The second product that was derived from the aerial photos is the canopy height model (CHM). To obtain the CHM, first, the digital surface model (DSM) is calculated using photogrammetry from overlapping air images. Then, the CHM is calculated by subtracting the DSM from the DTM. For this study, the Federal Forest Office (BFW) calculated the CHM.

Acquisition of the utilized digital terrain model (DTM) was performed by airborne laser scanning (ALS) in 2013 as part of the EU project INTERREG. The slope was calculated from the DTM using the Horn algorithm [29]. The Austrian Federal Forests provided all data.

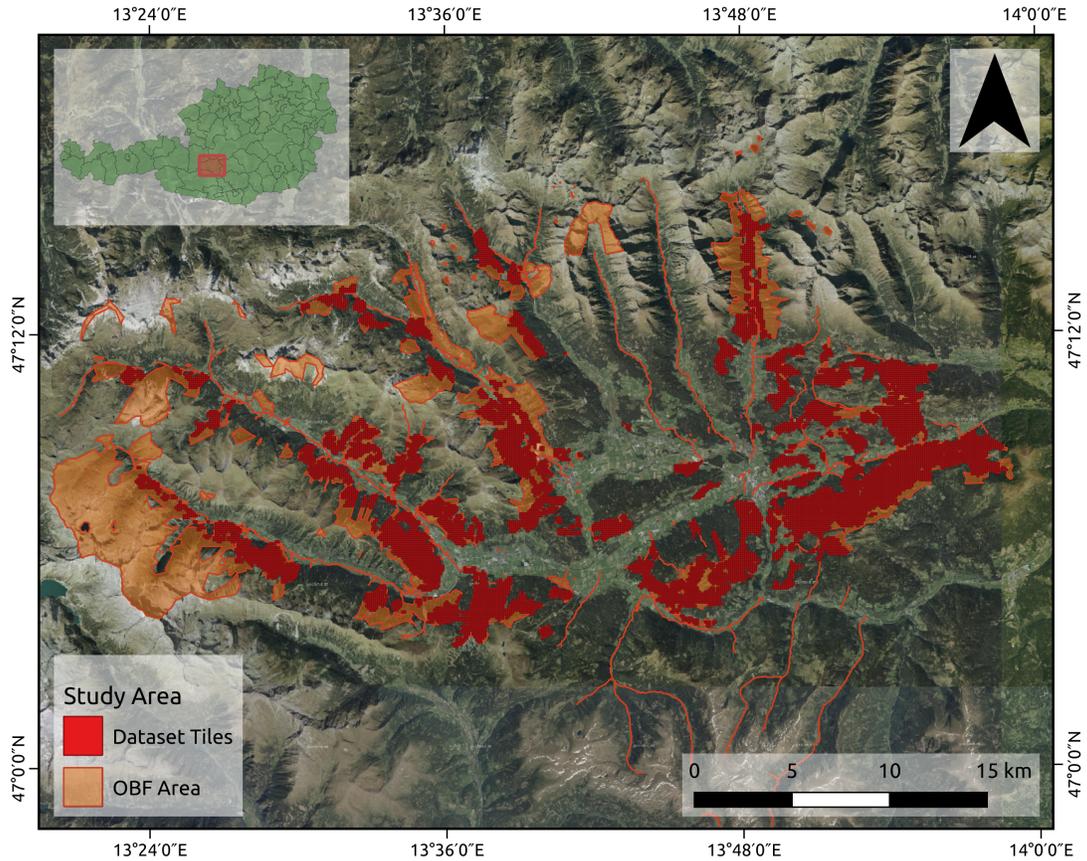


Figure 1: Study area of Lungau, Austria. The Austrian Federal Forest area is colored in orange, includes commercial and protective forest as well as non forest areas. Shown in red are the tiles that are used to create the data set that is used in this paper. The background image is a true color orthophoto from airborne photography [5] of the region in Austria that was used. The region in the main image is the same as the red square shown in the full map of Austria shown in the upper left.

Reference data collection was performed by forest engineers from May 2017 until November 2017 and subsequently digitalized by April 2018. This data acquisition was made as part of the forest management plan update by the Austrian Federal Forests. Forest engineers assess every forest stand by measuring the basal area and tree heights to derive the most important key figures such as tree species composition, yield class and stocking volume. Another part of the assessment is the planning of thinnings that need to be executed to ensure the optimal growth of the trees and maintain healthy forest stands with high-quality wood. At the Austrian Federal Forests thinnings are planned in three urgency levels, *i.e.* the forest stand:

- urgency 1 - needs thinning during the next 0 to 3 years
- urgency 2 - needs thinning during the next 3 to 10 years
- urgency 3 - can be thinned after 10 years or can be postponed until the next management plan

The ‘urgency level’ terminology is an internal annotation used by the Austrian Federal Forests as a practical approach to record the urgency for use by the forest manager. However, just the

first two (urgency 1 and urgency 2) are relevant, since by the time urgency 3 would be executed a new management plan will have been made. Therefore it is used very rarely and will not be considered in this study. The acquired data was populated into a Database (See Table 2), and the spatial information of the forest stands was drawn in a GIS. Subsequently, both data sources were spatially registered.

Table 2: Reference data classes, their definitions and occupied area.

Class	Definition	Area [ha]	Area [%]
thinning 1	Forest, thinning within 0-3 years	958.95	9.3
thinning 2	Forest, thinning within 3-10 years	1404.88	13.6
no thinning	Forest, no thinning	7245.62	70.3
other	Non forest (buildings, roads, water bodies)	696.19	6.8

3.2 Methods

The goal is to predict where and when thinnings are necessary. Deep convolutional neural networks (DCNNs) are used to classify based on pixels, i.e. semantic segmentation. This section describes the data preprocessing, outlines the experimental design, describes the training procedure, and presents the evaluation criteria applied.

3.2.1 Data preprocessing

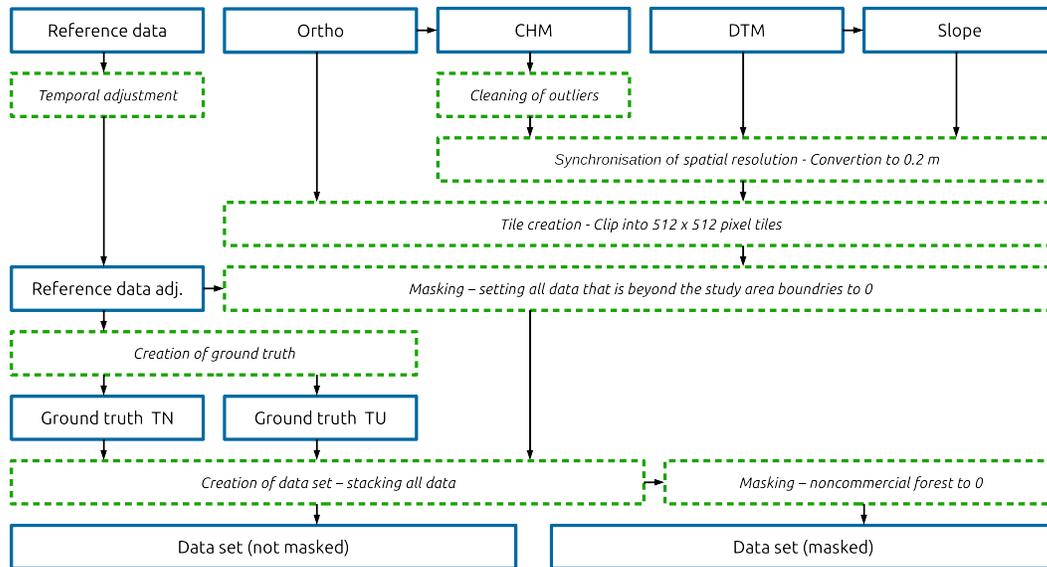


Figure 2: Data preprocessing workflow with all data manipulation steps (see text).

Due to different spatial and temporal resolutions (Table 1), the described data needed preprocessing before input into the DCNN. The workflow is shown in Figure 2 and described in this section in more detail.

The provided canopy height model (CHM) contained values ranging from -73.9 to 138. Since values below 0 are physically impossible and tree heights above 45 m are improbable in Austria, these values need adjustment. Therefore, values below 0 and over 40 m were highlighted, visually inspected and classified in GIS. The analysis showed that highlighted pixels higher than 47 m

and negative values were predominantly rocky steep slopes. Consequently, all pixels with a value beneath 0 and above 47 m were set to 0.

3.2.2 Synchronization of spatial resolution

Next, the spatial resolution of all input data was synchronized. As shown in Table 1, the orthophotos have a spatial resolution of 0.2 m as opposed to all other raster data having a spatial resolution of 1 m. There are benefits and drawbacks to using either as the standard spatial resolution. Using 0.2 m as the standard results in more detail in the four orthophoto layers. Greater detail can be beneficial for the deep learning algorithm to recognize individual tree crowns and thus help it determine the density of the forest. In contrast, the resolution of 1 m would provide a broader field of view when maintaining the same input image size to the DCNN. The resolution of 0.2 m was chosen as the standard for all input data as otherwise the anticipated information loss would be too high. This assumption is based on the fact that the average tree crown diameter for a first thinning is approximately 2.5 m. Consequently, a pixel size of 1 m would imply that one tree crown is represented by about 3 pixels in contrast to the 13 pixels with a pixel size of 0.2 m. The determination of an individual tree crown using 3 pixel resolution is much harder compared to a tree crown represented by 13 pixels. Consequently, all data with a spatial resolution of 1 m was converted using the GDAL library [16].

3.2.3 Tile size

The data needed to be clipped into square tiles as it is the input type of the neural networks used in this study. Accordingly, an image size of 512x512 pixels was a favorable choice due to the possibility of processing this image size with modern GPUs with a decent batch size while maintaining a large field of view. Ultimately the image size could be easily reduced by quartering the dataset without considerable effort, thereby increasing the batch size while processing accordingly. The image size of 512x512 pixels corresponds to 102.5m x 102.5m in reality. Since one adult tree crown (Norway spruce) has a diameter of 5-6 m, 300-400 adult tree crowns can be represented on one tile.

3.2.4 Tile creation

Only commercial forest is relevant for this study, since thinnings are only regularly planned and executed with commercial forests, as described in Section 2.1. Thus, a regular grid of 102.5m x 102.5m polygons was laid over the entire study area to create polygons that were then intersected with the reference data. Only polygons that contained over 25% commercial forest area were selected, while the remaining polygons were removed. Finally, the remaining polygons were used to clip all input raster data into 512 x 512 pixels tiles.

3.2.5 Reference data adjustment

Considering that the aerial imagery acquisition was recorded in September 2018 and the reference data from May to November 2017, the provided reference data needed to be adjusted. In the year between the two acquisitions, the study area was modified by planned and salvage logging.[65] This inconsistency was resolved using data from the Austrian Federal Forests database, where all cuttings are registered, to identify already carried out thinnings and salvage logging. Additionally, all commercial forest stands were visually inspected in GIS, and any noticeable errors were corrected.

3.2.6 Masking

Since the reference data is restricted to the study area and square tiles were used, part of the tile frequently contains no reference data (ground truth). Having parts of the tile with no ground truth information while orthophoto and DTM data is available would be misleading for the machine learning algorithm. Hence, any data outside the study area's boundaries was removed. A binary

mask from the reference data for every tile was created. In the resulting tiles, all information beyond the boundaries of the study area was set to 0.

Table 3: Allocation of the reference data to ground truths TN and TU as well as the proportions of pixels representing each class. TN is an abbreviation for thinning necessity and TU for thinning urgency. The definitions of the reference data classes are explained in Table 2.

Reference data Class	TN			TU		
	Class	Train[%]	Test[%]	Class	Train[%]	Test[%]
void	void	8.5	8.5	void	8.5	8.5
thinning 1	thinning	21.0	21.0	thinning 1	8.4	8.4
thinning 2				thinning 2	12.5	12.5
no thinning	no thinning	64.3	64.3	no thinning	64.3	64.3
other	other	6.2	6.2	other	6.2	6.2

Table 4: Allocation of the reference data to masked ground truths $masked\ TN$, $masked\ TU$ as well as the proportions of pixels representing each class. TN is an abbreviation for thinning necessity and TU for thinning urgency. The definitions of the reference data classes are explained in Table 2.

Reference data Class	masked TN			masked TU		
	Class	Train[%]	Test[%]	Class	Train[%]	Test[%]
void	void	44.2	44.2	void	44.2	44.2
thinning 1	thinning	20.4	20.4	thinning 1	8.4	8.4
thinning 2				thinning 2	12.0	12.0
no thinning	no thinning	35.4	35.4	no thinning	35.4	35.4

3.2.7 Design of ground truth label models

The ground truth was created from the reference data with its six unique classes, as shown in Table 3. "Void" represents the absence of information as described in Section 3.2.6. The other five classes are defined in Section 3.1.2, with three related to thinning. The first claim is that detecting forests that need thinning from remote sensing data is feasible. Since we are not interested in the acuteness of the thinnings for this objective, we summarized *thinning 1* and *thinning 2* into one class called *thinning*. This ground truth is called TN (thinning necessity), and it differentiates between the classes *void*, *thinning*, *no thinning* and *other* (Table 3). To address the second claim regarding the urgency of thinnings, we created another ground truth named TU (thinning urgency) (Table 3). TU was designed to differentiate between urgent thinning (*thinning 1*), not urgent thinning (*thinning 2*), *no thinning*, *other*, and *void*. Furthermore, we constructed two additional ground truths named $masked\ TN$ and $masked\ TU$ to focus exclusively on commercial forest (Table 4). The ground truths $masked\ TN$ and $masked\ TU$ expect to answer the claims as stated above for TN and TU , however all non commercial forest (mainly protective forest and roads) was excluded. Finally, all ground truth data was created by rasterizing the reference data (vector data) with the GDAL library into images of size 512x512x1.

3.2.8 Creation of data set

All the preprocessed input data and the generated ground truths were stacked into two data sets. Two separate data sets were needed since the input data corresponding to ground truths TN , and TU were masked differently to the input data of ground truths $masked\ TN$ and $masked\ TU$. Hence, the data sets stacked all preprocessed data into two hdf5 files. Each file consists of all tiles created inside the study area, as shown in Figure 1. Every tile is composed of the

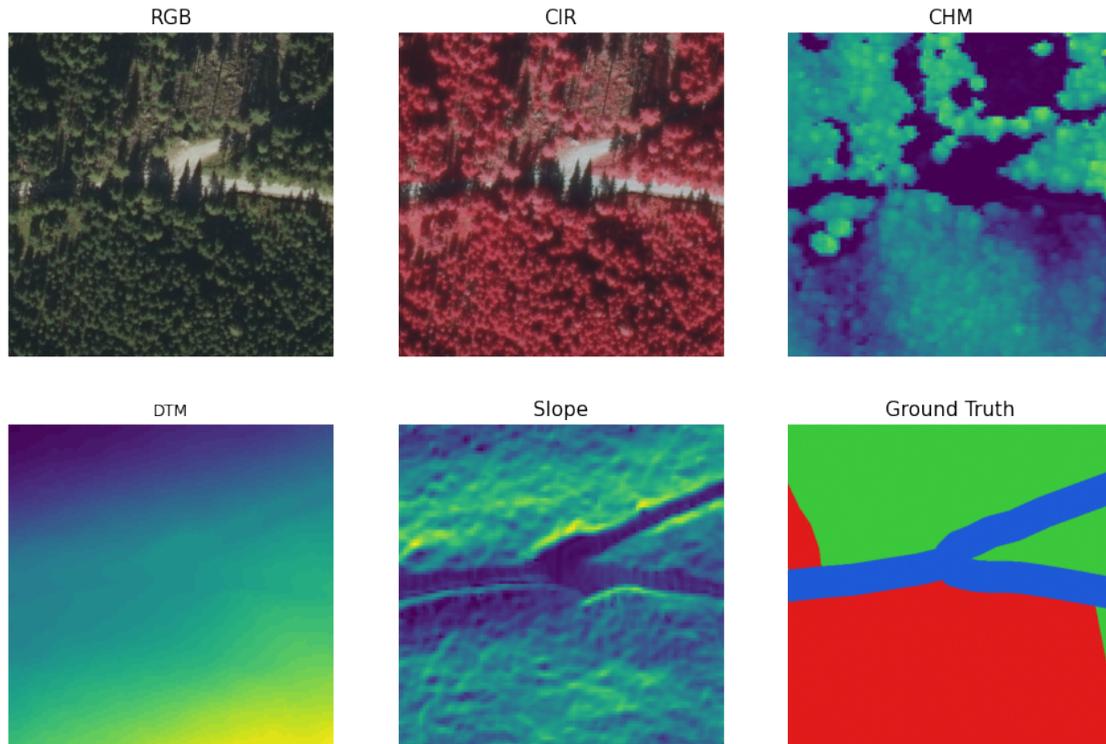


Figure 3: Input data tile as used in the final data set (non masked). RGB: true color orthophoto, CIR: the color infrared orthophoto, CHM: crown height model, DTM: digital terrain model, Slope: slope and Ground Truth: ground truth TN . Ground Truth classes are colored where green represents forest not to be thinned (class: forest), red represents forest thinning is necessary (class: thinning) and blue represents everything else (class: other)

orthophoto (RGB+NIR, 4 layers), the CHM (1 layer), the DTM (1 layer), the Slope (1 layer) and the Ground truths (3 layers), thus summing up to 10 layers (Figure 3). Thus, in total per data set, 10250 tiles were created, each of size $512 \times 512 \times 10$ pixels resulting in an array with the dimensions $10250 \times 10 \times 512 \times 512$. All input data were standardized by subtracting the mean and dividing by the standard deviation for each value of each input channel (1).

4 Results

4.1 Experimental design

The experimental design can be divided into three major parts and is illustrated in Figure 4. In the first part called *Architecture Selection*, a search for the best performing DCNN-architecture is performed. The results of the models predicting the necessity of thinning (*Thinning Necessity*) and the urgency of thinning (*Thinning Urgency*) are analyzed. Finally, an ablation study (*Ablation Study*) is performed to determine the impact of the various input data sources.

The exploration to find the best performing DCNN architecture was conducted exclusively on the data set with the ground truth TN (Table 3). The other ground truth type, TU used later to clarify if deep nets can determine thinning urgency, was not employed in searching for the best model. From the literature review reported in [53], the following three DCNNs were found as the most promising architectures for solving the semantic segmentation problem:

- UNet [52]

- FC-DenseNet [36]
- DeepLabv3+ [10]

The architecture utilized in this study is shown in Figure 8. The DCNN was modified to accept the tiles from the data set with the dimensions $512 \times 512 \times 7$ as input and output the predictions as a $512 \times 512 \times 1$ array. We search for the optimal architecture by evaluating the DCNNs on the data set (TN). The DCNN achieving the best result is then chosen for further optimization experiments. This optimization consists of manipulating individual parts of the architecture to further optimize the model’s performance by applying the Bayesian optimization method [57]. This method is an alternative to an exhaustive full grid search, which is computationally expensive. Instead, it more efficiently searches for optimal hyper-parameters. In practice, we sequentially altered critical structures in the DCNN architecture and always adopted the structure that provided the best results. All DCNNs are implemented in PyTorch, and all supplementary code was written in Python. All code is freely available at https://github.com/satlawa/edin_thinning_necessity. The experiments were performed on a system with Ubuntu 20.04 as the operating system equipped with an I7 6700K CPU, 32GB of RAM and an Nvidia RTX 3090.

After finding the best architecture, the final models were trained on the two data sets with all six ground truths (TN , TU , $masked\ TN$, $masked\ TU$). These final models were then evaluated on the test set to estimate the actual performance on unseen data. Subsequently, the results were analyzed to determine whether the models can satisfactorily fulfil the intended task. In the case of the ground truth TN and $masked\ TN$, we examined if the necessity of thinnings can be detected by a DCNN using solely remote sensing data. Furthermore, the ground truths TU and $masked\ TU$ were used to investigate if the urgency of thinnings can be determined in addition to detecting the necessity of thinnings.

Finally, an ablation study was performed by omitting different types of data from the input data set with ground truth TN . The primary purpose of this experiment was to determine the importance of the individual input data types, as shown in Table 1.

4.2 Training

The data set was randomly shuffled with 70% of the data used for training, 10% for validation and the remaining 20% for testing of all models. The test set is exclusively applied to the best

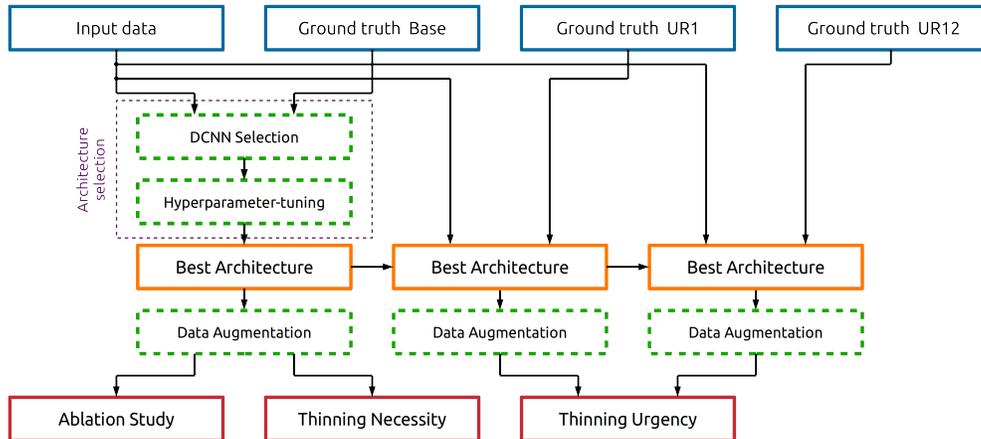


Figure 4: Model training workflow with all processing steps to obtain the final models. The first stage is the Architecture selection to find the best performing DCNN-architecture. The second stage is the training of the final models to answer the research objectives (*Thinning Necessity*, *Thinning Urgency*, *Ablation Study*).

performing model chosen by evaluation on the validation set.

After setting aside 20% of the data for the test set and 10% for validation, the remaining 70% of the data is randomly shuffled to train the model. All models were trained from scratch due to the dissimilarity of the input data compared to the data sets used on pre-trained models. For the initialization of the weights, the Kaiming uniform initialization was employed [22]. For all experiments, the Adam optimizer was used as in [39], with an initial learning rate for (1) UNet of 0.01, (2) FC-DenseNet of 0.003 and (3) DeepLabv3+ of 0.001. After every epoch, a decay rate of 0.995 was applied to the learning rate. Since loss functions play a decisive role in training models, choosing a suitable loss function is essential [35]. Here, the distribution of the classes is skewed. Therefore, the DICE loss was used as the loss function. The DICE loss is defined as $1 - F_1$ score (defined in Equ. 5). The batch size was chosen to be as large as possible, with the constraint being the memory of the GPU. Depending on the DCNN architecture, the batch size ranged between 16 and 72. However, due to long training times (up to 40 minutes for one epoch), only one 5-fold cross-validation was carried out on the model TN . All models were trained until convergence, which required at least for 50 epochs. Data augmentation in the form of horizontal flips was performed only on the training set and the previously chosen best-performing architecture. Data augmentation was restricted to horizontal flips due to the different environmental conditions on north and south facing slopes.

4.3 Evaluation Criteria

To evaluate the performance of the trained models thoroughly, five evaluation metrics were used. The first criterion is the overall accuracy (Acc) which is determined by dividing the number of all correctly classified pixels by the total number of pixels (Equation 1).

$$Acc = \frac{\sum_{i=1}^n (TP_i + TN_i)}{\sum_{i=1}^n p_i} \quad (1)$$

where n is the number of images; TP_i and TN_i are the number of correctly labeled pixels in image i ; p_i is the number of pixels in image i .

Although Acc is a very intuitive and thus common evaluation criteria, it is not very meaningful in cases where the distribution of classes is highly skewed, as is the case with the data set we utilize (Table 2). Therefore, the metrics precision (Equation 3), recall (Equation 2), IoU (Equation 4) and F_1 score (Equation 5) are used, for each class j :

$$recall_j = \frac{\sum_{i=1}^n TP_{ij}}{\sum_{i=1}^n (TP_{ij} + FN_{ij})} \quad (2)$$

$$precision_j = \frac{\sum_{i=1}^n TP_{ij}}{\sum_{i=1}^n (TP_{ij} + FP_{ij})} \quad (3)$$

$$IoU_j = \frac{\sum_{i=1}^n TP_{ij}}{\sum_{i=1}^n (TP_{ij} + FP_{ij} + FN_{ij})} \quad (4)$$

$$F_1^j = 2 * \frac{precision_j * recall_j}{precision_j + recall_j} \quad (5)$$

where n is the number of images; TP_{ij} is the number of pixels in image i , which are correctly predicted as class j ; FP_{ij} is the number of pixels in image i , which are incorrectly predicted as class j ; FN_{ij} is the number of pixels in image i , which are incorrectly predicted as any class other than class j .

Since this is a multi-class problem statement, we first calculate the proposed metrics per class and then determine the mean among all classes.

Despite providing all the above metrics for a holistic evaluation of the models, the F_1 score is used as the sole decisive evaluation score. Furthermore, class *void* carries essentially no information. However, it is still necessary as the data outside the study area boundaries of the commercial forest boundaries (in the case of masked ground truths) must be represented. Nonetheless, although the class *void* is required for training the DCNNs, it is omitted in the presentation of the results as it was nearly perfectly classified and carried no valuable information. Besides, its presence in the metrics would distort the results and indicate a better model performance than the particular model can achieve in reality. Finally, all evaluation on the test set was performed on the 512x512 data set.

4.4 Main Experimental Results

This section outlines and interprets the results of the performed experiments based on the previously described methods in Section 3.2. In particular, it presents the selection of the best performing model. In addition, it demonstrates the feasibility of using DCNN for classifying the urgency of thinnings as well as the ablation study.

4.4.1 Architecture Selection

DCNN Selection First, the performance of the three DCNN architectures is compared: UNet ($F_1 = 79.82\%$), FC-DenseNet ($F_1 = 72.62\%$) and DeepLabv3+ ($F_1 = 80.38\%$) on the TN dataset. DeepLabv3+ was chosen as the default network architecture and the hyper-parameter tuning was performed on this DCNN. More details on the selection process can be seen in [53].

Hyper-parameter tuning

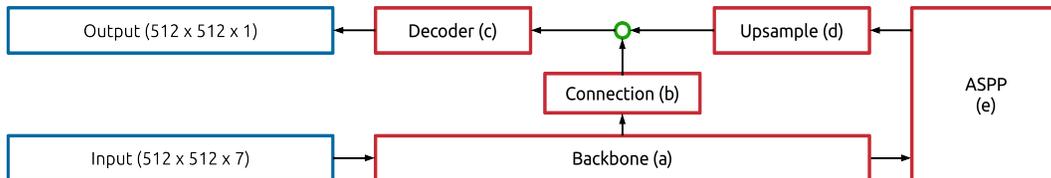


Figure 5: Modified parts of DeepLabv3+. Parts represent by letters.

Although the DeepLabv3+ architecture is already tuned on the PASCAL VOC 2012 data set, the data set used here is considerably divergent. Thus, hyper-parameters were optimized by applying the Bayesian optimization method [57]. The performance of the DCNN was explored by modifying five critical parts of the DeepLabv3+ as shown in Figure 5.

The optimization process began with exchanging the modified Xception architecture with the Resnet 101 architecture as the backbone module of the DCNN (part a). This module is responsible for encoding the features from the initial images until they are passed to the ASPP module. As Table 5 shows, replacing the backbone module to the Resnet 101 architecture results in an F_1 score of 81.26%, which is a 0.88% gain compared to the 80.38% achieved with the Xception architecture. Hence, Resnet 101 was chosen as the backbone for our DeepLabv3+ net for all further experiments.

Other variations were explored: to part b, ([1x1,48] convolution), part c ([1x1,256]*2 upsample), part d, ([3x3,256]*2 decoder) and part e (output stride 8) and augmentation with left-right flipping. Only the stride change and flipping gave an improvement, resulting in $F_1 = 83.01\%$. More details of the tuning can be found in [53]. The best performing hyper-parameters and DCNN-architecture were used to train all subsequent models. A detailed diagram of the final architecture is provided in the Appendix.

4.4.2 Thinning

After optimizing the DCNN-architecture on the data set, the main research objectives of this study were addressed. First, the possibility of detecting the need for thinning with the optimized DCNN was evaluated, exclusively with remote sensing data (Section 4.4.3). Then, the feasibility of detecting the need of thinning with urgency was assessed (Section 4.4.4).

4.4.3 Thinning necessity

Based on the network architecture from Section 4.4.1, the model was evaluated on the TN test set (Table 6). With a mean F_1 score of 82.23%, the model achieves a similar score to the 83.01% on the TN validation set. Thus, one can conclude that the chosen model is not overfitting to the validation set. Additionally, when focusing on the class-specific scores, the model performs best on predicting the class *no thinning*, whereas the scores of the other two classes are significantly lower.

When examining the confusion matrix in Table 7, one can gain further insight into the misclassifications of the model. Accordingly, the main mistakes are those between the classes *thinning* and *no thinning* as well as between *no thinning* and *other*, while misclassifications between the classes *thinning* and *other* are insignificant. These results match our knowledge about the classes. Examples of predictions are illustrated in Figure 6.

For instance, class *thinning* represents dense forest with a minimum top height of around 13 m and thus contrasts to class *other* that often represents no vegetation or shallow growing vegetation like grassland or mountain pines. In contrast, it makes sense to see the model misclassifying the class *no thinning* with both other classes as it embodies forest at all ages. For example, a very young forest has almost identical features compared to grassland, as presented in Figure 6 row 1. Therefore, it is challenging and sometimes impossible to distinguish between *no thinning* and *other* with only remote sensing data. Equally ambitious are some classification cases between the classes *thinning* and *no thinning*. In this situation, the model struggles to differentiate between edge cases of dense forest (Figure 6 rows 5 to 8). These errors might arise due to lacking information about the yield class of the forest. For instance, dense old forest on less vigorous sites might appear similar to a dense middle-aged forest on productive sites or vice versa (Figure 6 row 6). Likewise, young forest on very viable sites might already require thinning during the planning period whereas similar forest on less viable sites grows slower and should therefore not be planned for thinning.

Unquestionably, the model produces misclassifications, yet the results are excellent for most cases. For example, it correctly classifies the recently thinned forest as not to be thinned (Figure 6 row 2). Similarly, the model has no problems classifying correctly the cut forest and the more sparsely standing forest as *no thinning*, whereas the dense forest it accurately predicts as *thinning* (Figure 6 row 3 and 4).

As foresters are particularly interested in the differentiation between forest with and without

Table 5: Effect of backbone (part a) on the validation set performance (ground truth TN). Acc: overall accuracy, mIoU: mean intersection over union and F1: F_1 score

Backbone	Acc	Precision	Recall	mIoU	F1
Xception	85.83%	81.88%	79.47%	67.76%	80.38%
ResNet 101	86.81%	82.94%	79.85%	69.00%	81.26%

Table 6: Class scores and mean class scores on the test set (ground truth TN) with 5-fold cross validation. The model’s objective is to detect the necessity of thinning. Class definitions, *thinning*: thinning within 1-10 years, *no thinning*: no thinning necessary, *other*: non forest areas. std: standard deviation.

Score	thinning	no thinning	other	mean	std
Precision	77.08%	91.58%	79.69%	82.78%	0.78%
Recall	80.32%	90.92%	74.15%	81.79%	0.52%
IoU	64.81%	83.89%	62.34%	70.35%	0.44%
F1	78.64%	91.24%	76.80%	82.23%	0.31%

Table 7: Confusion matrix on test set (ground truth TN). The numbers represent classified pixels times 10^6 .

		prediction			Σ	
		Class	thinning	no thinning		other
reference	thinning		89	26	1	116
	no thinning		22	341	5	368
	other		1	7	24	32
	Σ		112	374	30	516

the necessity of thinning, another model was trained that discriminates just among the classes *thinning* and *no thinning*. For this experiment, the data set containing the ground truth *masked TN* was used, which just contains information about commercial forest while all other data was masked.

Table 8: Class scores and mean class scores on the test set (ground truth *masked TN*). The model’s objective is to detect the necessity of thinning restricted to the commercial forest.

Score	thinning	no thinning	mean
precision	76.98%	92.27%	84.63%
recall	85.84%	86.82%	86.33%
IoU	68.31%	80.93%	74.62%
F1	81.17%	89.46%	85.32%

As reported in Table 8, a mean F_1 score of 85.32% was obtained. Hence, by focusing on merely two classes, the mean F_1 score was raised by 2.85% compared to the data set without masking (Table 6). When concentrating on the class-specific scores, it can be seen that the mean F_1 score’s gain is due to the more accurate classification of the class *thinning*.

More detailed information of the misclassifications between the classes *thinning* and *no thinning* is provided in the confusion matrix in Table 9. What is striking is the lower number of categorized pixels due to the masking of non-commercial forest areas and the much higher number of false positives compared to the false negatives. This finding contrasts with the first model evaluated (Table 7), where the false positives and the false negatives were relatively balanced. Precisely this behavioral distinction can be observed in the examples of Figure 6. Particularly apparent is the difference in sensitivity in the fifth row of Figure 6, where the prediction of the non-masked data set (d) predicts *no thinning* on the entire tile while the prediction of the masked data set (f) classifies the entire commercial forest as in need of thinning.

4.4.4 Thinning urgency

The next research question addressed is whether it is possible to detect the need of thinnings and predict their urgency accurately. Compared to the model in subsection 4.4.3, this question adds another layer of complexity to the network since it has to assess the thinnings urgency. For

answering this question, the ground truth *masked TU* was used. The idea of ground truth *TU* is to predict thinnings and their urgency directly. Hence it differentiates between very urgent thinnings (*masked thinning 1*), less urgent thinnings (*masked thinning 2*) and no need for thinning *masked no thinning*. A detailed description of the ground truths is provided in the section 3.2.7. Only the masked version is considered here to focus on the commercial forest. Experiments with the full data can be seen in [53].

Constraining the problem by just taking into account commercial forest increased the performance slightly in the case of ground truth *masked TN* (subsection 4.4.3). Accordingly, the same strategy was applied on the model *masked TU*. However, the results show no significant overall improvement in the mean F_1 (Table 10) over model *TU*. Although the model *masked TU* achieves a higher F_1 score in class *thinning 1* (51.48%) compared to model *TU* (41.03%), all other class scores achieve worse performance. Since we are particularly interested in predicting the urgency of thinning, model *masked TU* provides a slightly better alternative than model *TU*. Thus one can conclude that the restricted model *masked TU* gives a small improvement in class *thinning 1* compared to the model *TU*.

Fortunately, perfect accuracy is not needed because the forestry management procedures can tolerate the slight inefficiency of postponing an urgent thinning until the next thinning review cycle.

4.4.5 Ablation study

An ablation study investigated the importance of the individual input features by training models with various combinations of input features removed. The factors explored were the inclusion or exclusion of these data planes: Orthophoto, CHM (Crown height model), DTM (Digital terrain model), Slope. The performance variations of the different models provide insight into what effect a specific input feature has on the model. See [53] for the full details of the experiments summarized below.

The results showed that the inclusion of the *DTM* has no or a slightly negative impact on the model’s performance. When examining the classification scores, it is primarily in the class *other* where the inclusion of *DTM* deteriorates the performance. The input feature *Slope* improves performance relative to models that do not contain it. Since *Slope* is calculated from the *DTM*, we expected that the DCNN could at least partially learn to extract some useful information out of the *DTM*. However, it seems the DCNN is not capable of deriving *Slope* out of the *DTM*. Possibly this is because the input features *Slope* and *DTM* provide only information about the terrain and not the forest itself.

Omitting the input feature *CHM* results in models having a lower score compared to the equivalent models with *CHM* included. Consequently, one can conclude that *CHM* contains some unique information since tree heights are essential criteria for assessing the necessity of thinning, as section 2.1 briefly outlines.

Omitting the input feature *Orthophotos* gives a significant decline in performance compared to the full model. It appears that *Orthophotos* contains the most valuable information of all input features. This result coincides with the finding that the input feature *Orthophotos* is also the most informative for human forest managers when assessing a forest.

Table 9: Confusion matrix on test set (ground truth *masked TN*). The numbers represent classified pixels times 10^6 .

		prediction		
		thinning	no thinning	Σ
ref.	thinning	97	16	113
	no thinning	29	191	220
Σ		126	207	333

5 Discussion

Timely planning and execution of thinnings are crucial for maintaining a healthy forest, but minimal research has been performed to derive the need for thinning directly from remote sensing data. Compared to detecting the necessity of thinnings, the estimation of inventory data [20, 6, 15] or the classification tree species [4, 14, 34, 3] has received much more attention. Accordingly, this study demonstrates the potential of predicting the necessity of thinning with state of the art deep learning architectures solely from high altitude remote sensing data.

Using multispectral orthophotos, canopy height data (CHM), a digital terrain model (DTM) and slope, and the reference data collected in the field by experts, two data sets were created to investigate the research question: can the necessity for thinning be reliably predicted. Using the DeepLabv3+ DCNN-architecture for semantic segmentation, we achieved an F_1 score of 82.2% and demonstrated that deep learning algorithms can be applied to accurately detect forests in need of thinning. Additionally, by employing the masked data set to restrict the analysis to exclusively the commercial forest, the approach improved the F_1 score to 85.3%.

From these results, one can conclude that the DCNN was able to extract critical information about the density of the forest from the remote sensing data, and thus to assess the need for thinning. In comparison, [33] used machine learning algorithms on satellite imagery to find forest stands with a thinning need within the next 10 years and achieved an overall accuracy of 64.1%. However, their achievements are based on Landsat TM satellite imagery with a lower spatial resolution of 30 m compared to 0.2 m and a stand wise classification of forest operations in contrast to semantic segmentation. Thus, the results are hardly comparable to this study. Another study from [63] used exclusively ALS data to predict the necessity of thinnings with an overall accuracy of 79%. There too, the necessity of thinnings was determined through a stand-wise classification of forest operations.

The necessity for thinning is determined mainly by two criteria: tree heights and the standing density (basal area) of a forest stand, and the tree heights are already part of the input data. Thus, it appears that the model was capable of inferring the crown density and not the actual basal area. Therefore, the crown density is sufficient for the target thinning type, which is crown thinning. For other thinning types, such as low thinning, the DCNN might struggle to produce comparably good results. The model might provide moderate performance since only suppressed and sub-dominant trees are removed in low thinning, which have no impact on the canopy.

Besides creating a model that detects forests in need of thinning, we additionally explored predicting the urgency of thinning. In this case, the trained models struggled to distinguish between urgent (51.48%) and not urgent (49.57%) thinning even though we focused just on commercially utilised forest. (see Table 10). The poorer performance is possibly due to inconsistency in the data and missing crucial information that is not contained in the input data. Consequently, adding additional data such as yield class or age could improve the results.

An ablation study explored the importance of the individual input features, showing that orthophotos contain the most critical information for the model that assesses thinning. Adding the CHM further improved the performance predicting thinning. On the other hand, the input features DTM and Slope seem not to contain any additional helpful information.

The proposed model can be seen as a cost-efficient and scalable solution for reducing the costs of creating new forest management plans. With this approach, the necessity of sending expensive

Table 10: Class scores and mean class scores on the test set (ground truth *masked TU*). The model’s objective is to detect the urgency of thinnings restricted to the commercial forest.

Score	thinning 1	thinning 2	no thinning	mean
precision	42.72%	50.32%	92.97%	62.00%
recall	64.74%	48.84%	83.19%	65.59%
IoU	34.66%	32.95%	78.27%	48.63%
F1	51.48%	49.57%	87.81%	62.95%

forestry personal into the field can be greatly reduced. Thus, with the reduced costs the frequency with which forests are planned can be shortened, thereby greatly improving the robustness of the stand against natural disturbances, increasing the quality of the wood, and optimising the revenue.

As stated earlier, the proposed model is specifically trained to detect the need of selective crown thinnings in spruce dominated forest stands in the study area. Consequently, further research could examine the feasibility of employing DCNNs for other thinning types, additional tree species and other geographical areas. The training of a comparable model for deciduous forests might be challenging because of the greater diversity of tree species and the more difficult task of segmenting deciduous tree crowns.

The performance of assigning the urgency of thinning was unsatisfying in this study. Hence, one could conduct further research on this objective by providing additional valuable data such as age or yield class. Another potential direction of research could be the prediction of the volume of harvested wood for sales planning. Applicable to both small forest owners with limited funds and also to large forest management of companies as a quick help, this approach provides the ability to infer critical information about the forest promptly.

Final Remarks Through the fusion of several types of remote sensing data and a suitable deepnet classifier, this paper presented a method to detect the necessity of thinning in spruce forests (c. 85% F1 score for identifying if forest stands need thinning). The approach was less accurate (c. 63% F1 score) when distinguishing between urgent and less urgent thinning needs. Although the resulting model needs further tuning for production use, the paper showed the potential of using remote sensing data to plan thinning cost-effectively.

Acknowledgements The research reported here was undertaken by author Satlawa while an employee of the Austrian Federal Forests Office (OBFAG). We are grateful for their support for the study and providing the data.

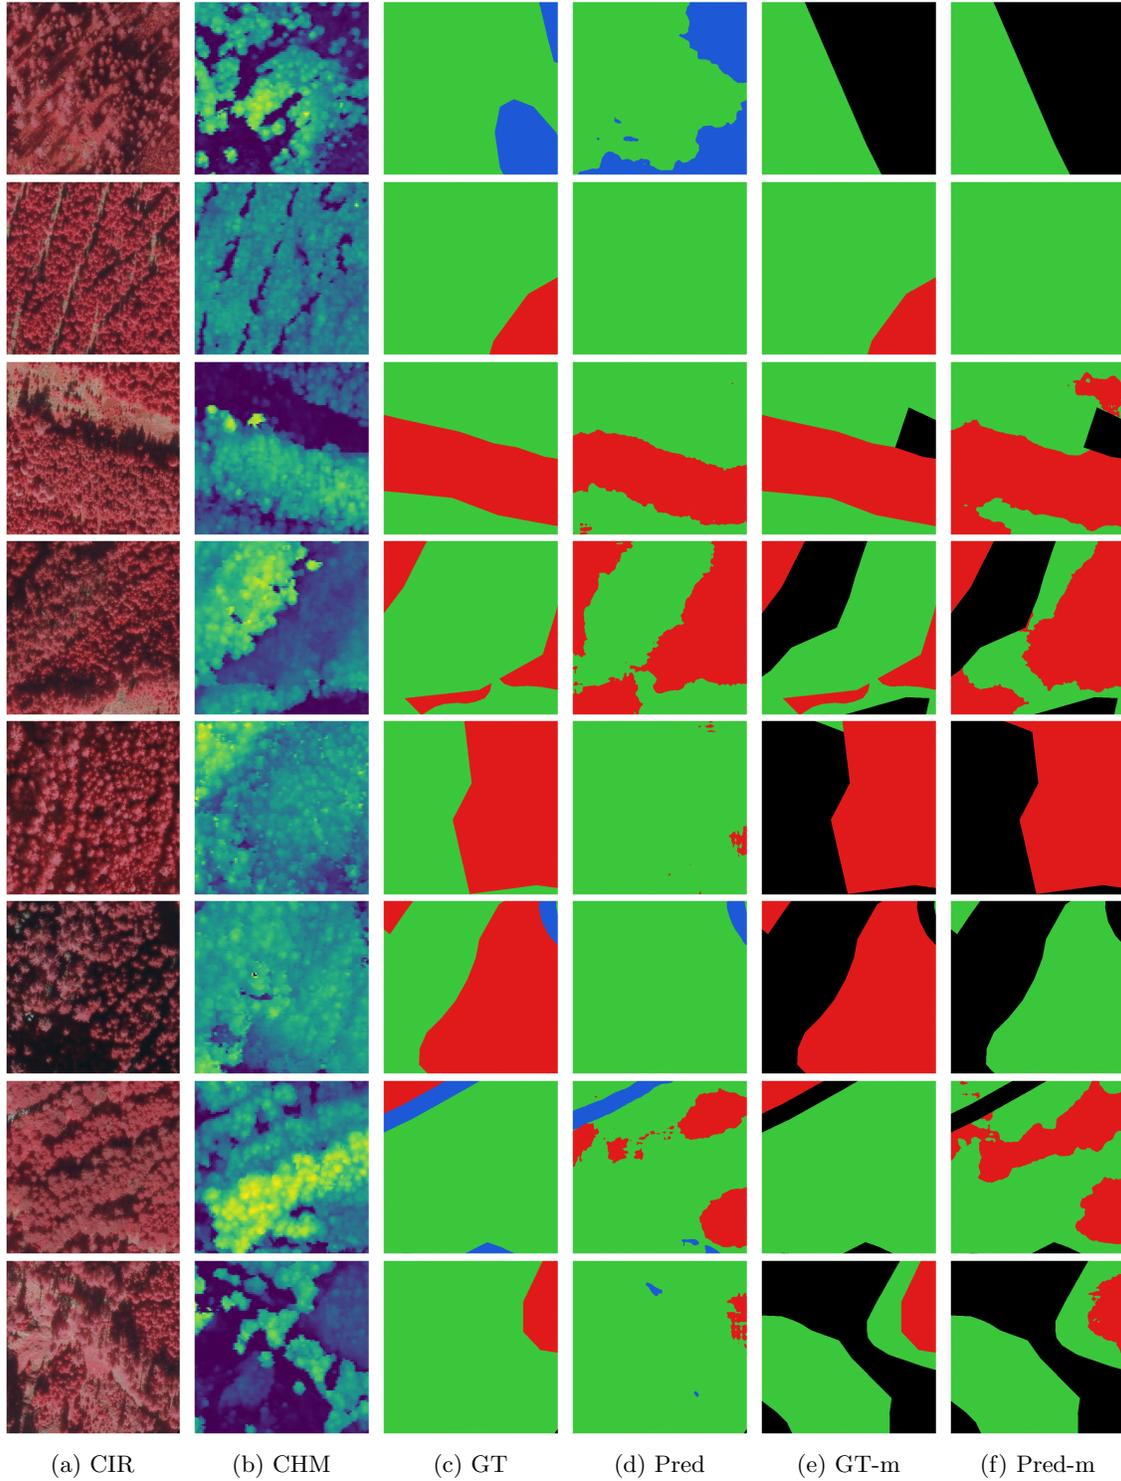


Figure 6: Examples of semantic segmentation on test set. Column (a) CIR are false color composites with near infrared, (b) CHM is the canopy height model, (c) GT is the ground truth of the Thinning Necessity set TN , (d) is the prediction of the model trained on the ground truth TN , (e) GT-m is the ground truth *masked* TN , (f) is the prediction of the model trained on the ground truth *masked* TN . In (b) the color palette illustrates low heights as dark (dark blue) and high heights as bright (yellow). The colored pixels in (c) to (f) represent the following classes, black: void, red: thinning, green: no thinning, blue: other.

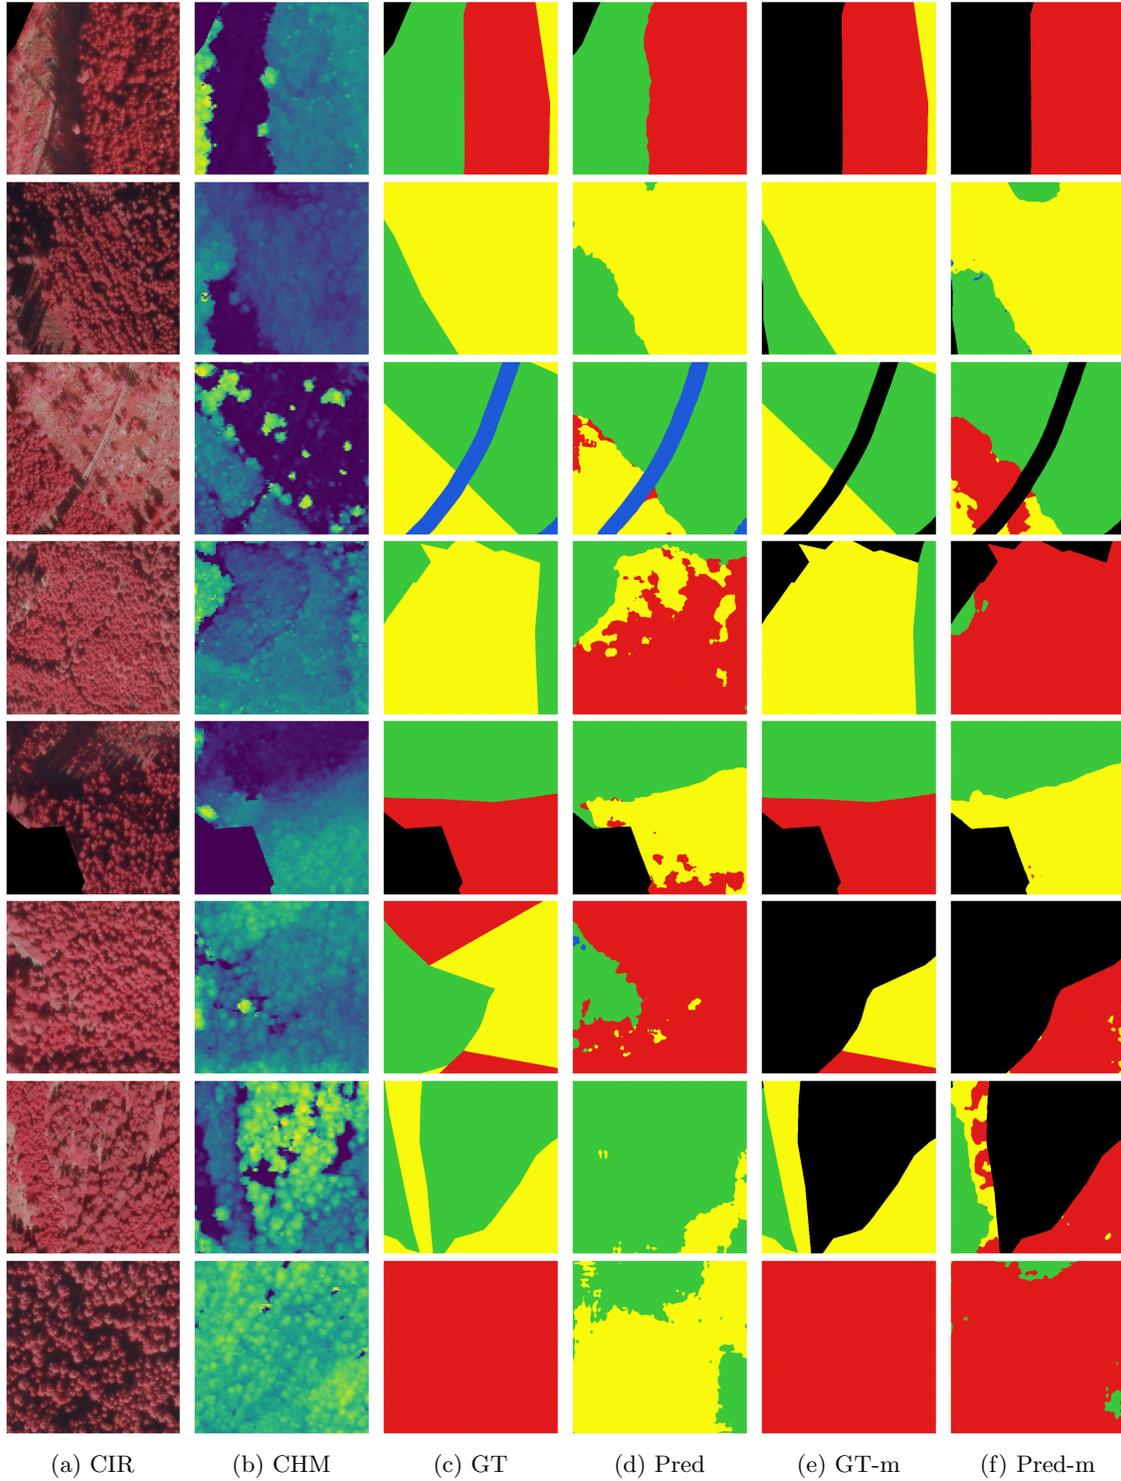


Figure 7: Examples of semantic segmentation on test set. Column (a) CIR are false color composites with near infrared, (b) CHM is the canopy height model, (c) GT is the ground truth Thinning Urgency set TU , (d) is the prediction of the model trained on the ground truth TU , (e) GT-m is the ground truth *masked* TU , (f) is the prediction of the model trained on the ground truth *masked* TU . In (b) the color palette illustrates low heights as dark (dark blue) and high heights as bright (yellow). The colored pixels in (c) to (f) represent the following classes, black: void, red: thinning 1, yellow = thinning 2, green: no thinning, blue: other.

Appendix: Final DCNN architecture

Figure 8 shows the final DCNN-architecture tuned on the data set (ground truth Thinning Necessity set TN). This architecture was used to train all models to resolve the research objectives "thinning necessity" (Subsection 4.4.3), "thinning urgency" (Subsection 4.4.4) and the ablation study (section 4.4.5).

The diagram (Figure 8) illustrates the overall structure of the network. However, batch norm layers and RELU activation functions are omitted due to space constraints. A batch norm layer follows every convolution as well as atrous convolution. The RELU activation functions follow every Resnet-101 convolution block (fine dotted violet lines), otherwise every convolution and atrous convolution. The numbers on the right side of the ResNet-101 blocks express how many times the block is repeated. The information flow into the Connection happens after the first Resnet-101 block is repeated three times. The PyTorch implementation is provided at https://github.com/satlawaw/edin_thinning_necessity.

Abbreviations The following abbreviations are used in this manuscript:

ALS	airborne laser scanning
ASPP	Atrous Spatial Pyramid Pooling
CHM	canopy height model
DCNN	deep convolutional neural network
DTM	digital terrain model
FCN	fully convolutional network
GIS	Geographical Information System
LD	Linear dichroism
NIR	near infrared
RGB	Red, Green, Blue

References

- [1] P Abetz. Biologische produktionsmodelle als entscheidungshilfen im waldbau. *Forstarchiv*, 41:5–9, 1970.
- [2] Thomas Eugene Avery. *Forester's guide to aerial photo interpretation*. Number 308 in Agriculture Handbook. US Department of Agriculture, Forest Service, 1966.
- [3] Arvid Axelsson, Eva Lindberg, Heather Reese, and Håkan Olsson. Tree species classification using Sentinel-2 imagery and Bayesian inference. *International Journal of Applied Earth Observation and Geoinformation*, 100(April 2020):102318, 2021.
- [4] Laurel Ballanti, Leonhard Blesius, Ellen Hines, and Bill Kruse. Tree species classification using hyperspectral imagery: A comparison of two classifiers. *Remote Sensing*, 8(6):1–18, 2016.
- [5] basemap. Orthofoto tilecache of austria, published by geoland.at, 2021.
- [6] Jonas Bohlin, Inka Bohlin, Jonas Jonzén, and Mats Nilsson. Mapping forest attributes using data from stereophotogrammetry of aerial images and field data from the national forest inventory. *Silva Fennica*, 51(2):1–18, 2017.
- [7] A. D. Cameron. Importance of early selective thinning in the development of long-term stand stability and improved log quality: A review. *Forestry*, 75(1):25–35, 2002.
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014.

- [9] Liang Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2018.
- [10] Liang Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11211 LNCS:833–851, 2018.
- [11] Ziyi Chen, Cheng Wang, Jonathan Li, Wentao Fan, Jixiang Du, and Bineng Zhong. Adaboost-like End-to-End multiple lightweight U-nets for road extraction from optical remote sensing images. *International Journal of Applied Earth Observation and Geoinformation*, 100:102341, 2021.
- [12] S. Daume and D. Robertson. A heuristic approach to modelling thinnings. *Silva Fennica*, 34(3):237–249, 2000.
- [13] Foivos I. Diakogiannis, François Waldner, Peter Caccetta, and Chen Wu. ResUNet-a: A deep learning framework for semantic segmentation of remotely sensed data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 162(March 2019):94–114, 2020.
- [14] Geoffrey A. Fricker, Jonathan D. Ventura, Jeffrey A. Wolf, Malcolm P. North, Frank W. Davis, and Janet Franklin. A convolutional neural network classifier identifies tree species in mixed-conifer forest from hyperspectral imagery. *Remote Sensing*, 11(19), 2019.
- [15] Selina Ganz, Yannek Käber, and Petra Adler. Measuring tree height with remote sensing—a comparison of photogrammetric and LiDAR data with different field measurements. *Forests*, 10(8), 2019.
- [16] GDAL/OGR contributors. *GDAL/OGR Geospatial Data Abstraction software Library*. Open Source Geospatial Foundation, 2022.
- [17] Pedram Ghamisi, Naoto Yokoya, Jun Li, Wenzhi Liao, Sicong Liu, Javier Plaza, Behnood Rasti, and Antonio Plaza. Advances in Hyperspectral Image and Signal Processing: A Comprehensive Overview of the State of the Art. *IEEE Geoscience and Remote Sensing Magazine*, 5(4):37–78, 2017.
- [18] M. Gruber, M. Ponticelli, R. Ladstädter, and A. Wiechert. Ultracam Eagle, Details and Insight. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXIX-B1(September):15–19, 2012.
- [19] Arto Haara and Kari T Korhonen. Toimenpide-ehdotusten simulointi laskennallisesti ajantasaistetusta kuvioaineistosta. *Metsätieteen aikakauskirja*, 2004(2):157–173, 2004.
- [20] Eelis Halme, Petri Pellikka, and Matti Möttöus. Utility of hyperspectral compared to multi-spectral remote sensing data in estimating forest biomass and structure variables in Finnish boreal forest. *International Journal of Applied Earth Observation and Geoinformation*, 83(April):101942, 2019.
- [21] Kaiming He and Jian Sun. Convolutional neural networks at constrained time cost. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 07-12-June-2015(3):5353–5360, 2015.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *Proceedings of the IEEE International Conference on Computer Vision*, 2015 Inter:1026–1034, 2015.

- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1904–1916, sep 2015.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2016-Decem:770–778, 2016.
- [25] S. Hein, S. Herbstritt, and U. Kohnle. Auswirkung der Z-Baum-Auslesedurchforstung auf Wachstum, Sortenertrag und Wertleistung im Europäischen Fichten-Stammzahlversuch (Picea Abies [L.] Karst.) in Südwestdeutschland. *Allgemeine Forst- und Jagdzeitung*, 179(10-11):192–201, 2008.
- [26] Per Holgén, Ulf Söderberg, and Björn Hånell. Diameter increment in Picea abies shelterwood stands in northern Sweden. *Scandinavian Journal of Forest Research*, 18(2):163–167, 2003.
- [27] Markus Holopainen, Mikko Vastaranta, and Juha Hyypä. Outlook for the next generation’s precision forestry in Finland. *Forests*, 5(7):1682–1694, 2014.
- [28] Dariia Hordiuk, Ievgenii Oliinyk, Volodymyr Hnatushenko, and Kostiantyn Maksymov. Semantic segmentation for ships detection from satellite imagery. In *2019 IEEE 39th International Conference on Electronics and Nanotechnology (ELNANO)*, pages 454–457. IEEE, 2019.
- [29] Berthold K.P. Horn. Hill Shading and the Reflectance Map. *Proceedings of the IEEE*, 69(1):14–47, 1981.
- [30] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-Janua:2261–2269, 2017.
- [31] Jari Hynynen, Anssi Ahtikoski, Juha Siitonen, Risto Sievänen, and Jari Liski. Applying the MOTTI simulator to analyse the effects of alternative management schedules on timber and non-timber production. *Forest Ecology and Management*, 207(1-2 SPEC. ISS.):5–18, 2005.
- [32] P Hyvönen, J Heinonen, and A Haara. Detection of forest management operations using Bitemporal aerial photographs. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences - ISPRS Archives*, 38(1999):309–313, 2010.
- [33] Pekka Hyvönen. Kuvioittaisten puustotunnusten ja toimen- pide-ehdotusten estimointi kälähimmän naapurin menetel- mällä Landsat TM -satelliittikuvan, vanhan inventointitiedon ja kuviotason tukiaineiston avulla [Estimation of stand characteristics and forest management op. *Metsätieteen aikakauskirja*, 2(2002):363–379, 2002.
- [34] Markus Immitzer, Martin Neuwirth, Sebastian Böck, Harald Brenner, Francesco Vuolo, and Clement Atzberger. Optimal input features for tree species classification in Central Europe based on multi-temporal Sentinel-2 data. *Remote Sensing*, 11(22), 2019.
- [35] Shruti Jadon. A survey of loss functions for semantic segmentation. *2020 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology, CIBCB 2020*, 2020.
- [36] Simon Jegou, Michal Drozdal, David Vazquez, Adriana Romero, and Yoshua Bengio. The One Hundred Layers Tiramisu: Fully Convolutional DenseNets for Semantic Segmentation. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 2017-July:1175–1183, 2017.

- [37] A. Juodvalkis, L. Kairiukstis, and Rimvydas Vasiliauskas. Effects of thinning on growth of six tree species in north-temperate forests of Lithuania. *European Journal of Forest Research*, 124(3):187–192, 2005.
- [38] Annika Kangas, Rasmus Astrup, Johannes Breidenbach, Jonas Fridman, Terje Gobakken, Kari T. Korhonen, Matti Maltamo, Mats Nilsson, Thomas Nord-Larsen, Erik Næsset, and Håkan Olsson. Remote sensing and forest inventories in Nordic countries—roadmap for the future. *Scandinavian Journal of Forest Research*, 33(4):397–412, 2018.
- [39] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pages 1–13, 2015.
- [40] Cheng-Chien Liu, Yu-Cheng Zhang, Pei-Yin Chen, Chien-Chih Lai, Yi-Hsin Chen, Ji-Hong Cheng, and Ming-Hsun Ko. Clouds classification from sentinel-2 imagery with deep residual learning and semantic image segmentation. *Remote Sensing*, 11:119, 01 2019.
- [41] Man Liu, Bolin Fu, Shuyu Xie, Hongchang He, Feiwu Lan, Yuyang Li, Peiqing Lou, and Donglin Fan. Comparison of multi-source satellite images for classifying marsh vegetation using DeepLabV3 Plus deep learning algorithm. *Ecological Indicators*, 125:107562, 2021.
- [42] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 07-12-June, pages 3431–3440. IEEE Computer Society, oct 2015.
- [43] Strütt M. Betriebswirtschaftliche modelluntersuchungen zu z-baum orientierten produktionsstrategien in der fichtenwirtschaft. *Mitteilungen der Forstlichen Versuchs- und Forschungsanstalt Baden-Württemberg*, 156:221, 1991.
- [44] Lei Ma, Yu Liu, Xueliang Zhang, Yuanxin Ye, Gaofei Yin, and Brian Alan Johnson. Deep learning in remote sensing applications: A meta-analysis and review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 152(April):166–177, 2019.
- [45] Elspeth Macdonald, Barry Gardiner, and William Mason. The effects of transformation of even-aged stands to continuous cover forestry on conifer log quality and wood properties in the UK. *Forestry*, 83(1):1–16, 2010.
- [46] R F MacKenzie. Silviculture and management in relation to risk of windthrow in Northern Ireland. *Irish Forestry*, 33:29–38, 1976.
- [47] Steen Magnussen, Thomas Nord-Larsen, and Torben Riis-Nielsen. Lidar supported estimators of wood volume and aboveground biomass from the Danish national forest inventory (2012–2016). *Remote Sensing of Environment*, 211(January):146–153, 2018.
- [48] Stephen J. Mitchell. Stem growth responses in Douglas-fir and sitka spruce following thinning: Implications for assessing wind-firmness. *Forest Ecology and Management*, 135(1-3):105–114, 2000.
- [49] Aire Olesk, Kaupo Voormansik, Mirjam Põhjala, and Mart Noorma. Forest change detection from Sentinel-1 and ALOS-2 satellite images. *Proceedings of the 2015 IEEE 5th Asia-Pacific Conference on Synthetic Aperture Radar, APSAR 2015*, pages 522–527, 2015.
- [50] P Persson. Windthrow in forests-its causes and the effect of forestry measures [sweden, scots pine, norway spruce]. *Rapporter och Uppsatser-Skogshoegskolan, Institutionen foer Skogsproduktion (Sweden)*, 1975.
- [51] J. Pollanschutz. Erfahrung aus der schneebruchkatastrophe 1979. *Allegem. Forstztg.*, 91(5):123–125, 1980.

- [52] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9351, pages 234–241. Springer Verlag, 2015.
- [53] P. Satlaw. Detecting the necessity of thinnings with deep learning. Master’s thesis, University of Edinburgh, School of Informatics, <https://homepages.inf.ed.ac.uk/rbf/DISSERTATIONS/Satlaw21.pdf>, 2021.
- [54] Jean Philippe Schütz, Michael Götz, Willi Schmid, and Daniel Mandallaz. Vulnerability of spruce (*Picea abies*) and beech (*Fagus sylvatica*) forest stands to storms and consequences for silviculture. *European Journal of Forest Research*, 125(3):291–302, 2006.
- [55] Hein Sebastian, Herbstritt Stephan, and Ulrich Kohnle. Auswirkung der z-baumauslesedurchforstung auf wachstum, sortenertrag und wertleistung im europäischen fichtenstammzahlversuch (*picea abies* [l.] karst.) in südwestdeutschland. *Allgemeine Forst und Jagdzeitung*, 179, 01 2007.
- [56] Marian Slodicak, Jiri Novak, and Jens Peter Skovsgaard. Wood production, litter fall and humus accumulation in a Czech thinning experiment in Norway spruce (*Picea abies* (L.) Karst.). *Forest Ecology and Management*, 209(1-2):157–166, 2005.
- [57] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical Bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, 25:2960–2968, 2012.
- [58] H Spellmann and M Schmidt. Massen-, sorten- und wertertrag der fichte in abhangigkeit von der bestandesbehandlung. *Forst und Holz*, 58(13/14):412–419, 2003.
- [59] Graham Stirling, Barry Gardiner, Elspeth Macdonald, Shaun Mochan, and Thomas Connolly. Stem straightness in Sitka spruce in South Scotland. Technical report, Forest Research, 2000.
- [60] Andrei Stoian, Vincent Poulain, Jordi Inglada, Victor Poughon, and Dawa Derksen. Land cover maps production with high resolution satellite image time series and convolutional neural networks: Adaptations and limits for operational systems. *Remote Sensing*, 11(17):1–26, 2019.
- [61] Erik Valinger and Jonas Fridman. Models to assess the risk of snow and wind damage in pine, spruce, and birch forests in Sweden. *Environmental Management*, 24(2):209–217, 1999.
- [62] Erik Valinger, Mikael Ottosson Lövenius, Ulf Johansson, Jonas Fridman, Svante Claeson, and Åsa Gustafsson. Analys av riskfaktorer efter stormen Gudrun. Technical report, Swedish Forest Agency - Skogsstyrelsen, 2006.
- [63] Mikko Vastaranta, Markus Holopainen, Xiaowei Yu, Juha Hyypä, Hannu Hyypä, and Risto Viitala. Predicting stand-thinning maturity from airborne laser scanning data. *Scandinavian Journal of Forest Research*, 26(2):187–196, 2011.
- [64] Michele Volpi and Devis Tuia. Deep multi-task learning for a geographically-regularized semantic segmentation of aerial images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 144(July):48–60, 2018.
- [65] Wikipedia. Salvage logging, 2023.
- [66] Xiaowei Yu, Juha Hyypä, Harri Kaartinen, and Matti Maltamo. Automatic detection of harvested trees and determination of forest growth using airborne laser scanning. *Remote Sensing of Environment*, 90(4):451–462, 2004.
- [67] Kai Yue, Lei Yang, Ruirui Li, Wei Hu, Fan Zhang, and Wei Li. TreeUNet: Adaptive Tree convolutional neural networks for subdecimeter aerial image segmentation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 156(April 2018):1–13, 2019.

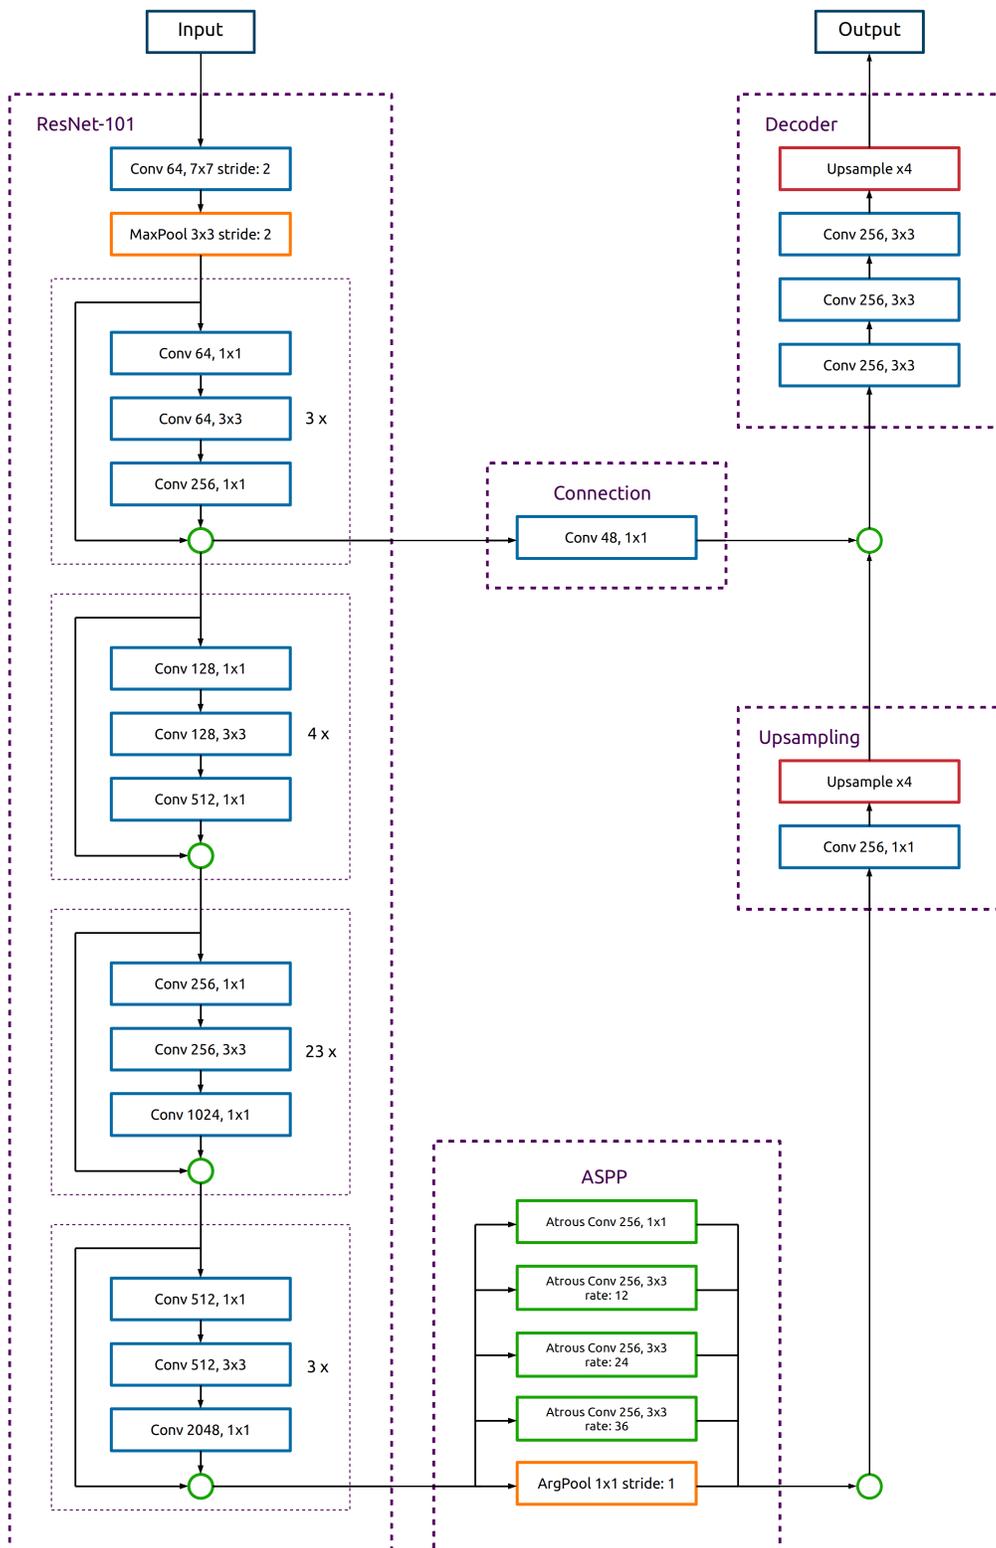


Figure 8: Diagram of DeepLabv3+ with ResNet-101 as backbone [10]. This network architecture is the final deep convolutional neural network used to train all models to detect the necessity and the urgency of thinnings. Green circles: concatenations.