

Norwegian University of Life Sciences

Master's Thesis 201930 ECTSFaculty of Environmental Sciences and Natural Resource Management

Mapping forest key habitats using machine learning and remote sensing data

Kartlegging av nøkkelbiotoper i skog ved bruk av maskinlæring og fjernmålingsdata

Jon Hidle Pedersen Geomatics Faculty of Science and Technology

1 Abstract

The Norwegian forest habitat inventory (Miljøregistrering i skog, MiS), is a system used to classify different forest key habitats for use in forestry planning. The purpose of the system is to localize and map habitats in managed forests that are of particular importance for biodiversity. The purpose is to take better care of the biodiversity in the forest of landowners, something that is important for sustainable forest management.

Today the work with identifying MiS habitats is largely based on fieldwork and does only to a small degree use information from remote sensing data. Creating a probability map of the location of MiS habitats can be a method of utilizing remote sensing data. One reason to make this map is giving field workers aid when determining what area of the forest to survey. Work can then be done more efficiently and targeted. This is because field workers spend their time searching the forest for the many different habitats. Therefore, a probability map of the different habitats can reduce the search time, by excluding certain areas where the possibility of finding MiS habitats is low.

Remote sensing data, like aerial photographs, can be a useful tool when creating a probability map of MiS locations. Aerial imagery is in studies shown to be a useful tool for finding objects in the forest. For example, aerial photographs have been used to classify different tree species. Utilising more of the information in remote sensing data can be a smart way to improve MiS mapping. Convolutional neural networks, a form of machine learning, have been successfully applied to a wide array of tasks including image recognition. Making a model learn how to recognize features and find patterns in data on their own, is the idea behind machine learning. A model can differentiate, for example, a dog image from a cat image, by using features it has learned by looking at many other cat and dog images. The idea behind this thesis was that if machine learning can be used successfully on many different image recognition tasks, could it also be used to find MiS habitats in aerial images? Could a useful probability map of MiS locations be created?

In an attempt to find MiS habitats in the study area, a convolutional neural network was created. This network used small sections of aerial images, image tiles, as input and output a prediction on what habitats it thinks is in the tile. Data from the fieldwork was used to label each image tile. The model was trained and validated on a training dataset before being tested on an independent test dataset. Several different models were tested, both multilabel and binary. In the end, it proved difficult to create a model that gave good and meaningful results. The best model managed to classify 61% of the test images correctly. This is a score too low to make this a method useful for field workers. It created a probability map of the different MiS habitats. This map was not useful because large parts were classified incorrectly, i.e. as MiS habitats when being something else, compared to the survey done in the field work.

How could the results of this work been improved? Model performance could have improved with more knowledge about input data. Knowing what features characterises the different MiS habitats may have helped tailor the model to be more suitable for this specific problem. Another advantage is knowing what parts of the data is valuable, and how to separate it from noise. Trying different approaches would also change the performance of the model; instead of using the image tiles, using bigger or smaller part of the aerial photographs would effect how the model behaves. The conclusion is that this thesis used a method that did not work with the procedure that was used. It is belived that more domain knowledge and a different approach may have given different and better results.

2 Samandrag

MiS (miljøregistrering i skog) er eit system som blir brukt til å klassifisera ulike nøkkelbiotopar til bruk i skogbruksplanar. Formålet med systemet er å lokalisera og kartlegga habitat i skog som er særleg viktige for biologisk mangfald. Målet er å ta betre vare på mangfaldet i skogen til kvar grunneigar, noko som er viktig for berekraftig skogforvaltning.

I dag er arbeidet med å identifisera MiS-habitat i stor grad basert på feltarbeid, og brukar i liten grad informasjon frå fjernmåling. Å laga eit sannynligheitskart over plasseringa til MiS-habitata kan vera ein måte å bruka slike fjernmålingsdata på. Ein grunn til å laga dette kartet er å gi feltarbeidarar hjelp til å avgjerda kva område av skogen som skal saumfarast. Arbeidet kan då gjerast meir effektivt og målretta. Dette er fordi ein feltarbeidar elles brukar mykje tid på å leita i skogen etter dei mange MiS-habitata. Difor kan eit sannynligheitskart som viser MiS-habitata redusera tida ein brukar på leita ved å utelukka visse område der det er lite sannsynleg å finna MiS-habitat.

Fjernmålingsdata, som flyfoto, kan vera eit nyttig verktøy for å laga eit sannsynligheitskart over MiS-habitat. Ulike studiar har vist at flyfoto kan vera eit godt reiskap for å finna ulike objekt i skog. For eksempel har flyfoto blitt brukt til å klassifisera ulike treslag. Å utnytta meir av informasjonen frå fjernmålingsdata kan vera ein smart måte å forbetra MiS-kartlegginga på. Konvolusjonelle nevrale nettverk, ei form for maskinlæring, har blitt brukt i ei rekke oppgåver, inkludert bildegjenkjenning. Å laga ein modell som på eiga hand lærer å gjenkjenna eigenskapar og å finna mønster i data, er ideen bak maskinlæring. Ein slik modell kan skilja for eksempel eit hundebilde frå eit kattebilde ved å bruka mønster han har lært seg ved å sjå mange bilde av kattar og hundar. Tanken bak denne oppgåva er at viss maskinlæring kan brukast med suksess på mange ulike bildegjenkjenningsoppgåver, kan det då brukast til å finna MiS-habitat i flyfoto? Kan det lagast brukande sannsynligheitskart for MiS-habitat?

I eit forsøk på å finna MiS-habitat i eit vald studieområde blei det laga ein maskinlæringsmodell. Dette nettverket brukte små utsnitt av flyfoto, bildefliser, som input og gav ut eit mål på kva habitat det trudde var i bildeflisa. Data frå feltarbeid vart brukt til å kategorisera bildefliser. Modellen vart trena og validert på eit treningsdatasett før han blei testa på eit separat testdatasett. Fleire ulike modellar blei testa, både multilabel og binære. Til slutt viste det seg vanskeleg å laga ein modell som ga gode og meiningsfulle resultat. Den beste modellen klarte berre å klassifisera 61% av testbilda rett. Dette er eit for dårleg resultat til å bli ein nyttig metode for ein feltarbeidar. Når ein brukte dette til å laga sannsynligheitskart for Mis-habitata vart det ikkje brukande fordi store deler av bildeflisene vart klassifisert feil, dvs som MiS-habitat utan å vera det, i forhold til registreringane frå feltarbeidet.

Korleis kunne resultata av dette arbeidet ha blitt betre? Modellytinga kunne blitt betre med meir kunnskap om inndata. Å forstå kva eigenskapar som karakteriserer dei ulike MiS-habitata kunne ha blitt brukt til å skreddarsy modellen slik at han er meir eigna for å løysa dette spesifikke problemet. Ein annan fordel er å vita kva delar av datamaterialet som er verdifulle og korleis skilje dette frå støy. Andre tilnærmingar ville også ha endra resultatet til modellen; i staden for å bruka bildefliser, vil større eller mindre utsnitt av flyfotoa påverke korleis modellen presterer. Konklusjonen er at denne oppgåva brukte ein metode som ikkje virka med den framgangsmåten som blei brukt. Likevel er det truleg at meir fagkunnskap og ei endra tilnærming kan gi eit anna og betre resultat.

Contents

1	Abstract	i
2	Samandrag	iii
3	Introduction3.1MiS, miljøregisteringer i skog3.2Hyperspectral imaging3.3Artificial neural networks3.4Convolutional neural networks3.5Objectives	1 2 3 4 5 6
4	Materials	7
	 4.1 Study area 4.2 Field data collection 4.3 Hyperspectral data 	7 7 8
5	Methods 5.1 Processing raw data 5.2 Labeling of data 5.3 Convolutional neural network 5.3.1 Preprocessing of images 5.4 Predication	 9 10 10 10 11 12
6	5.5 Evaluating model performance	12
U	6.1Evaluating objective one6.2Evaluating objective two6.3Evaluating objective three6.4Comparison of the objectives	13 16 18 21
7	Discussion 7.1 Problems encountered	22 23
8	Conclusion	23
9	Acknowledgments	24
10	Bibliography	25

3 Introduction

The Norwegian forest habitat inventory (Miljøregistrering i skog, MiS), is a system used to classify different forest key habitats for use in forestry planning [1], [2]. The purpose of the system is to localize and map habitats in managed forests that are of particular importance for biodiversity. This information is used when making forestry plans. The purpose is to improve the environmental considerations of each landowner, making it a basis for sustainable forest management. Examples, where these maps are used, is when planning new forest roads, deciding what forest areas that are scheduled for cutting, and when and how they are cut. Today the work with identifying MiS habitats is largely based on fieldwork and does only to a small degree use information from remote sensing data.

Creating a map of the location of MiS habitats can be a method to utilize remote sensing data. Such a map could provide field workers with information determining what area of the forest to survey. Thus, field work can then be done more efficiently and targeted, reducing the time a field worker spends searching the forest for the different habitats. Therefore, a probability map of the different habitats can be used to exclude certain areas where the possibility of finding key habitats is low and hence cut down the search time. This means that the field worker can survey a forest in a shorter amount of time, since his search area is reduced. Cost saving is an additional benefit of streamlining the search. reducing man-hours spent in the forest. Another advantage of using this kind of tool is the reduction of human error. Field workers ability to be objective can be affected by many different factors like the weather, fatigue, and experience. For example, rain can negatively affect the mood of a field worker, making the person rush through an area. By reducing the search area in the forest the objectiveness of the survey may improve and areas of interest is not left out or forgotten. For example, a study by Eriksen et al. [3] found that individual observers often classified the same area differently. From this, one could argue that a tool like a MiS probability map could reduce the effect of the individual observer and improve objectiveness.

Remote sensing data, like aerial photographs, can be a useful tool when creating a probability map of MiS locations. A method of aerial photography often used is hyperspectral imaging. This method collects more spectral information than ordinary images, making identification of objects in the image easier. Examples by [4] and [5] shows that remote sensing data can be used to find standing dead wood. Standing dead wood is one of the 12 MiS categories, see figure 1 for the different categories. Another MiS habitat thought to be easy to find is recently burned forest [6]. It is understandable that these two habitats are possible to identify. Looking at aerial images it is easy to spot brown dead trees or the big scars burnt areas leave in the terrain. On the other hand, habitats like lying dead wood can be harder to identify [7]. They may not be immediately visible from photographs, but indicators like specific habitats or structures where lying dead wood often can be found can give information that can be utilized to identify such areas [7]. For example, it can be logical to think that there is a correlation between standing dead wood and laying dead wood [8]. Another example is ground vegetation in the forest. It is like laying dead wood

difficult to see from the air but could be found looking at other indicators. Aerial photographs can be used to classify different tree species [9]. This can make it possible to identify specific habitats based on the dominant tree type in the area.

The development in machine learning algorithms has accelerated greatly since the world gained interest in 2012 [10]. Making a model learn how to recognize features and patterns in data, by itself, is the idea behind machine learning. Manually creating good general features by hand is difficult, and takes a long time. Justification for using machine learning on this problem compared to other methods is the ability to extract features and information from images. A machine learning model can differentiate, for example, a dog image from a cat image, by using features it has learned by looking at many other cat and dog images [11]. An example of a machine learning model's ability to classify images is the ImageNet [12] classification challenge. Here modern models can distinguish between hundreds of different classes, like dogs and cats but also car and football. Thus, the successfull use of machine learning on ImageNet to distinguish hundreds of classes, provide a foundation to use such methods to identify the 12 different MiS habitats. In theory, a model looking at many aerial images of MiS habitats could learn how to find MiS habitats in a forest. A study done by the Norwegian computing center found that using a tool like machine learning could increase the classification accuracy when classifying tree species compared to other techniques [13].

Machine learning was used by researchers to find and label four different tree species in forest imagery [14]. This resulted in a segmentation where nearly 4/5 of the trees were correctly classified.

A third example is how sickness in trees was found using machine learning [15]. Researchers used the spectral signatures of leaves to differentiate between sick and healthy trees with high precision.

3.1 MiS, miljøregisteringer i skog

MiS (miljøregisteringer i skog) [1] is a scientifically based registration procedure. The methodology was developed based on a project where the goal was to find areas in the forest particularly important for vulnerable species. Based on the results of this study, a registration method was developed for recording different environmental habitats in the forests.

The MiS registrations are based on field registrations based on instructions where criteria for design and content are determined. The registration is standardized but may be adapted to the local area. The registrations take place primarily in old forest and in areas where the forest is being cultivated. Twelve main habitat types were selected for the inventory, and they were delimited in the field according to specified criteria, either by using densities of habitat elements or by using vegetation or topographic criteria. The goal is to find and register habitats that are of particular importance for biodiversity [2]. Example of these habitats is old trees, dead standing trees, and ravines.

After the forest is mapped, the habitats are ranked and prioritized by value in a process where local considerations are taken into account. The process is set up to involve forest owners, and the most valuable habitats enter the forest owner's forest management plan. Municipal and regional authorities can use the information in their planning activities. The selected MiS registrations are the main basis for establishing the key habitats that will be managed in accordance with the guidelines in the forest industry. MiS is a tool used to maintain sustainable forestry.

The MiS registration method was created in collaboration with many different research institutions and have since 2004 been one of the standards used in Norwegian forest management planning. MiS data is open to the public.



Figure 1: The 12 MiS habitats. Image source: [2]

Twelve important habitats for biodiversity in managed forests

1 Snags (insects, fungi, birds, bats)

2 Logs (fungi, insects, bryophytes) 3 Trees with nutrient-rich bark (lichens, bryophytes)

4 Trees with pendant lichens (insects, spiders, mites, lichens)

5 Late successions of deciduous trees (insects, fungi, birds)

6 Old trees (insects, spiders, mites, lichens)

7 Hollow deciduous trees (insects, bats, birds)
8 Recently burned forest (fungi, insects, plants)
9 Luxuriant ground vegetation (plants, insects, fungi, snails)
10 Rock walls (bryophytes, lichens)
11 Clay ravines (lichen, bryophytes)
12 Stream gorges (bryophytes, lichens)

3.2 Hyperspectral imaging

Hyperspectral imaging [16], is a method that collects image information from across a broader, and higher spectral resolution, range of the electromagnetic spectrum than ordinary imagery. Each image represents a narrow wavelength range of the electromagnetic spectrum, also known as a spectral band. These 'images' are combined to form a three-dimensional (x,y,λ) hyperspectral image cube for processing and analysis, where x and y represent two spatial dimensions of the scene, and λ represents the spectral dimension (comprising a range of wavelengths). This is shown visually in figure 2).

Whereas ordinary cameras captures the color of visible light in mostly three broad bands, spectral imaging divides the spectrum into multiple contiguous narrow bands, often more than 100 bands. Hyperspectral images contains more spectral information compared to ordinary images, witch makes finding objects, identifying materials, or detecting processes easier.

Certain objects leave unique 'patterns' in the electromagnetic spectrum. Such patterns are known as spectral signatures. These patterns enable identification of the materials that make up a pictured object. It is these kinds of patterns this thesis hopes to find and use to distinguish between the 12 MiS habitats.



3.3 Artificial neural networks

Artificial neural networks (ANN) [18] is a system that processes data vaguely inspired by biological neural networks that make up an animal brain. A neural network is simply put a collection of neurons and connections between them. A 'neuron' is a function with a bunch of inputs and one output. Its task is to take a number as input, perform a function on it and output the result. See figure 3.



Neurons are connected to each other in layers. Neurons in a layer connect to the next and previous layers in a network and are not interconnected in the same layer. Data travels through the network in one direction, from the inputs of the first layer to the outputs of the last. A network with many layers is called a deep neural network. Each neuron connection has only one parameter called weight. These weights tell the neuron to respond more to one input and less to another. Weights are adjusted when training the network, that's how the network learns. An ANN is used in this thesis as a tool to predict the MiS classes.

3.4 Convolutional neural networks

In deep learning, a convolutional neural network (CNN, or ConvNet) is a class of deep neural networks, commonly applied to analyzing images [20]. A CNN consists of one or more convolutional layers and a deep neural network. The convolutional layer consists of a set of learnable filters. Every filter is a small 'window' looking at the image. During the training of a filter, we slide (convolve) each filter across the input image and compute dot products between the weights of the filter and the image. As we slide the filter over the image it produces an activation map that gives the responses of that filter. When training a layer, the filters will learn to recognize simple features like, for example, edges. Stacking multiple convolutional layers give the filters the ability to recognize more complex features by combining simpler features from earlier layers. See figure 4.



Convolutional layers with their weights can be reused in new models by using a process called transfer learning [21]. With transfer learning, instead of starting the learning process from scratch, the model already knows patterns that have been learned when solving a different problem [22]. This way one can use previous learnings from an already trained model, to build upon and improve a new model. A pretrained model is a model that was trained on a large dataset to solve a problem like the one that we want to solve. One example of a pretrained model is DenseNet that is trained on a large collection of images called ImageNet [12]. ImageNet contains hundreds of image classes like plants, animals, furniture, and vehicles.

Since filters are trained on input images, CNN's use relatively little preprocessing compared to other image classification algorithms. This means that the network learns the filters that in traditional algorithms were made by hand. This independence from prior knowledge and human effort in feature design is a major advantage. MiS habitats can therefore, in theory, be found without domain knowledge.

3.5 Objectives

The goal of this thesis is to evaluate the potential to provide a probability map of the location of forest key habitats using remote sensing data. This thesis is going to use a CNN to try to classify different MiS habitats using only images. The high grade of information in a hyperspectral image, combined with a CNN's ability to automatically extract features from images, is the foundation of why this method may succeed in classifying the different MiS habitats.

Three different objectives are explored in this thesis. The main difference between the objectives is what data is available to the model and how this data is processed. Each objective is made to explore different machine learning models ability to solve the MiS classification problem. The map made from each model could give different information based on what the models value. Below are the three objectives:

- 1. Find out if a model manages to differentiate between the 12 different MiS habitats in a multilabel classification task.
- 2. Find out if a model manages to differentiate between MiS areas and non-MiS areas in a binary classification.
- 3. Find out if the model improves by reducing the variation in the dataset.

4 Materials

4.1 Study area

The study area in this project is located near Biri, in Gjøvik municipality, Norway. (6756907, 585477, EUREF89 UTM-zone 32N). The size of the area is approximately 16 km^2 . The area consisting mostly of boreal forest and some farmland and is dominated by managed spruce forest. The area was split into two sub-areas used to train and test the model. See figure 5.

4.2 Field data collection

A field inventory was carried out in the area in 2017 to map MiS areas according to the standard field protocol. There was found nine of the twelve MiS categories, in an area covering about 6 % of the total study domain. Habitat distribution can be seen in table 1. Figure 5 shows the resulting MiS areas overlaid on the study area.

Habitat	count	percentage
Stående død ved	36.0	30.3
Liggende død ved	54.0	45.4
Rikbarkstrær	3.0	2.5
Trær med hengelav	1.0	0.8
Eldre lauvsuksesjoner	5.0	4.2
Gamle trær	4.0	3.4
Hule lauvtrær	0.0	0.0
Brannflater	0.0	0.0
Rik bakkevegetasjon	5.0	4.2
Bergvegger	6.0	5.0
Leirraviner	0.0	0.0
Bekkekløfter	5.0	4.2

Table 1: MiS habitats found inside the study area.

Figure 5: The study area in Biri shown with the location of MiS key habitats. Outlined are the two separate areas used to train and test the models.



4.3 Hyperspectral data

The hyperspectral data were collected in the summer of 2018 using two HySpex [23] hyperspectral sensors from Norsk Elektro Optikk. One image sensor operates in the short-wave infrared spectrum (SWIR) and the other sensor operates in the visible and near-infrared spectrum (VNIR). Together the sensors cover a spectral range from 400nm to 2500nm. Flight data and specification is listed in table 2.

Table 2: Flight data and	d sensor specification
Contractor	Terratec AS
Altitude above ground	900 meters
Area covered	$15.9 \mathrm{~km2}$
Datum	WGS84 UTM 32
SWIR sensor	HySpex SWIR-384
SWIR bands	288
SWIR spectral range	400nm to 1000 nm
VNIR sensor	HySpex VNIR-1800
VNIR bands	186
VNIR spectral range	$1000\mathrm{nm}$ to $2500\mathrm{nm}$

The sensors collected data at 900m above the terrain giving the SWIR sensor a pixel size, GSD [24], of 0.7 meters and the VNIR sensor a pixel size of 0.3 meters. Initial calibration and processing of the image data were done by TerraTec. The data was delivered in several overlapping image stripes.

5 Methods

The Python programming language with the libraries Keras [25] and GDAL [26] was the main tool used in this thesis to process the raw data into the finished result. All code used the thesis can be found on GitHub: https://github.com/JonHidlePedersen/Master-thesis-Jon-Hidle-Pedersen.

In an attempt to classify the different MiS habitats, a CNN was created. This network used small sections of aerial images, image tiles, as input and output a prediction on what habitats the model thought was in the tile. Data from the fieldwork, marking the location of the different habitats, was used to label each image tile. The model was trained and validated on a training dataset before being tested on a separate test dataset. Several different models were tested, both multilabel and binary. Using the predictions made on the test dataset it created a map showing where it was likely to find the different MiS habitats. The generalised version of the workflow is shown in figure 6.



Figure 6: Flowchart showing the steps to transform the raw observations into a MiS probability map.

5.1 Processing raw data

To get a manageable image size out of the big hyperspectral image stripes, it was decided to split each aerial stripe into tiles of 50x50 pixels. This gave each tile a real-world area of 250 m^2 when using the 0.3-meter GSD of VNIR. An area of 250 m^2 is often used when doing fieldwork. Processing the large aerial photographs into small tiles were done using GDAL's retile function. It was planned to merge the images from the SWIR and VNIR sensor to create a single image containing all the bands from both sensors. This was not done because the computation took too long to finish. Therefore, tiles were only made from the VNIR sensor. The 13 raw VNIR stripes could not be merged together to one big image. Some methods were tried with varying success and in the end the raw image stripes were used to create tiles. This caused some overlapping tiles and the MiS probability maps, therefore, have stripe artifacts. Having merged the image stripes would have removed these artifacts.

5.2 Labeling of data

Next step in the process was to label each tile using MiS data collected during the fieldwork. The location of each tile was compared to the location of the different MiS habitats. If a tile was inside a MiS habitat it was given the label corresponding to the habitat. Tiles not contained in any of the habitats was labeled as a new 'empty' class. This meant that each tile could be labeled as one, or multiple, of 13 (12 MiS + 1 empty) different classes. Tiles were also labeled as either test- or training tiles to ensure an unbiased evaluation of the final model performance, see figure 5. All tiles were also labeled using a forest mask. The labels were stored in a text file containing the tile name and the different labels. Labeling resulted in a total of about 2400 MiS tiles in the training dataset and about 1200 tiles in the test set. A model can thus be trained and evaluated on the study area using the existing MiS locations as a ground truth. Not all of the twelve classes were found in the dataset. The class distribution is shown in table 1. Balancing of the classes was done to ensure that the model did not favour any class more than the others. The twelve MiS classes were balanced using a weight matrix. The model was evaluated on the test set containing 1100 MiS tiles and 24000 non-MiS tiles. Note that the test set is predominantly non-MiS tiles. It is an imbalanced classification problem. This can make the model appear more capable than it is in reality because it can categorise everything as non-MiS and still get a good accuracy score. We are on the other hand interested in how good it is at finding MiS locations/tiles. Therefore, a few evaluation metrics is considered to get a better picture of the model performance.

5.3 Convolutional neural network

The image tiles could now be used in a CNN model created in Keras, and objective one could be tried. Importing the hyperspectral multilabel data into the model was more challenging than previously thought because some helper functions in Keras could not be used. The helper tools were made for ordinary image types, like a standard RGB image, and did not work on multiband hyperspectral images. Therefore, feeding images into the model, splitting data into trainingvalidation datasets, and data argumentation had to be done by custom functions. Making the custom functions took some time and did not end up as advanced and feature-rich as the functions provided by Keras. Nevertheless, the custom functions worked and gave the model image data and labels.

Two types of models were tried out, multilabel and binary. The design of each model was made using an iterative process of trial and error. First training and validating the current model, noting its performance, then tweaking the model before repeating the training process. During training, the model was saved during its best each training step, epoc, to ensure that the model was not overfitting. Only the model from the best training step was evaluated. What modifications should be made to the model was mainly decided by looking at the loss and accuracy plots made during the training process. MiS habitats are not mutually exclusive which meant that the first models were made to handle multilabel input and output. The performance of the model using the criteria from objective one is studied in the results chapter.

In objective two binary classification was tried. Binary classification was defined as MiS habitat or no MiS habitat. All 12 MiS classes were gathered into one. It was hoped that the binary model would score better than the multilabel model since it must only take a binary choice. The balancing of the two classes was done by reducing the number of non MiS tiles the model was trained on. This meant the model was trained on an equal amount of MiS and non MiS tiles. By reducing the 12 classes into one, the dataset is also made bigger on a class basis. Here the idea to use transfer learning was tried out. The pretrained model used in this thesis was DenseNet201 because it could handle the resolution of the 50x50 pixel tiles. One problem with using DenseNet was that it was pretrained on a large collection of images called ImageNet [12]. Imagenet consists mostly of three band images, not hyperspectral images with hundreds of bands. DenseNet therefore only accepts image types similar to what it was trained on, not hyperspectral images. This meant that the hyperspectral image tiles had to be changed from their hyperspectral form to a three-band image. The converted three-band tiles could be imported into the pretrained model. This was done using the built inn Keras helper functions. The performance of the model using the criteria from objective two is studied in the results chapter.

The last objective was to test a model on a reduced dataset. A forest mask was applied to the dataset witch limited the image tiles to only contain forest. This was done in an attempt to improve the accuracy score of the models by removing noisy images containing houses and roads. After all the model is trying to find MiS habitats, and filtering out disturbing objects could improve performance. The same model and procedure as in objective two were tried on the reduced dataset. The performance of the model using the criteria from objective three is studied in the results chapter.

5.3.1 Preprocessing of images

Converting the hyperspectral tiles to three band images was done using two methods. The first method tried was to stack each band side by side three bands deep creating a long narrow image. This method did not alter any information from the original image, but in testing, it did not perform well. The second method was to use principal component analysis. PCA is a technique used to emphasize variation and bring out strong patterns in a dataset [27]. PCA was used to reduce the image dimensionality down to the three most important PCA-components. The three first components accounted for about 80%, 10% and 3% of the total variation in the images. The remaining 183 components totaling 7% of the variation in the images where discarded. PCA processing does reduce the information in the image tile, compared to the original tile, but using this conversion method worked better than the band stacking. The tiles used in training and testing were converted and stored as three-band images.

5.4 Predication

Finally three MiS probability maps of the test area were made using the model predictions. When a image tile from the test area was given to a model a prediction was returned. This value was stored together with the bounding box of the image tile. Using Python and the GDAL library this information was used to create a polygon. Performing this prosses over all the image tiles creates a collection of polygons with individual MiS predictions. This collection is then stored as the resulting MiS probability map. QGIS [28], a geographic information system (GIS), was used to make the maps. The resulting map looks like a heatmap showing the possible occurrence of MiS habitats.

5.5 Evaluating model performance

To evaluate the performance of the model the following evaluation metrics where used: accuracy score, precision, recall, f1-score, confusion matrix and Cohen's kappa.

The accuracy score explains what percentage of tiles are correctly classified. Precision and recall help us understand how well our model is capturing information and how much it is leaving out. Precision tells us from all the test examples that were assigned a label, how many actually were supposed to be categorized with that label. Recall tells us, from all the test examples that should have had the label assigned, how many were actually assigned the label. In other words, precision is a measure of result relevancy while recall is a measure of how many truly relevant results are returned. A precision and recall score close to one is wanted, but a low score shows that the model is not very good at classifying tiles. The precision-recall curve shows the trade-off between precision and recall of different thresholds. A high area under the curve represents both high recall and high precision. High scores for both show that the classifier is returning accurate results (high precision), as well as returning a majority of all positive results (high recall). A system with high recall but low precision returns many results, but most of its predicted labels are incorrect when compared to the training labels. A system with high precision but low recall is just the opposite, returning very few results, but most of its predicted labels are correct when compared to the training labels. An ideal system with high precision and high recall will return many results, with all results labeled correctly [29].

The f1-score is the harmonic mean of precision and recall. Cohen's kappa coefficient is an evaluation metric which measures how much homogeneity, or consensus, there is in the ratings given by multiple judges. It is generally thought to be a more robust measure than simple percent agreement calculation, as Cohen's kappa takes into account the possibility of the event occurring by chance. A kappa score is wanted as high as possible. The confusion matrix is a table

describing the performance of a classification model by looking at how it correctly and incorrectly classifies data.

6 Result

Models made using criteria specified in the three objectives were evaluated on the test dataset. All results from each of the tree objectives are presented below. A comparison of the results is presented at the end of the chapter. Note that the test set is predominantly non-MiS tiles, it is an imbalanced classification problem.

6.1 Evaluating objective one

The multilabel models provided classification results with low accuracy. The best model provided an accuracy score of about 5 % on the test data, it gave no meaningful results. By looking at the multilabel model's evaluation metrics it seems that the model overfitted by ignoring several of the MiS classes, and only cared about the most frequent MiS class and the 'empty' class. To mitigating this problem a balancing of the classes and weighted labels were tested. This did not work and the multilabel approach was rejected in favor of binary classification.

The accuracy score on the test set was 5 % correctly classified tiles. The classification report lists the precision, recall, and f1-score of the two classes the model uses. See table 3. A high score for the MiS class is wanted but a low score shows that the model is not very good at finding MiS habitats. See figure 7.

1				
	precision	recall	f1-score	support
Standing dead wood	0.03	0.07	0.04	530
Lying dead wood	0.02	0.65	0.04	471
Trees with nutrient-rich bark	0.00	0.00	0.00	0
Trees with pendant lichens	0.00	0.00	0.00	83
Late successions of deciduous trees	0.00	0.00	0.00	0
Old trees	0.00	0.00	0.00	0
Hollow deciduous trees	0.00	0.00	0.00	0
Recently burned forest	0.00	0.00	0.00	0
Luxuriant ground vegetation	0.00	0.00	0.00	63
Rock walls	0.00	0.00	0.00	0
Clay ravines	0.00	0.00	0.00	0
Stream gorges	0.00	0.00	0.00	394
Not MIS	0.96	0.13	0.22	24104

Table 3: Evaluation metrics for the multilabel classification. Classification report:

Figure 7: The precision-recall curve to the multilabel model. It shows how bad the model is at finding the correct classes, its average precision being only 0.09. Ideally the graph should be a horizontal line with a high area under the curve. This graph shows the opposite with having a low area under the curve. Notice how fast the line declines!



Cohen's kappa gave a result close to zero indicating that the model gives no useful information. The confusion matrix shows that the multilabel model favors only two classes. This may be explained by the fact that many of the image tiles contain one of the two classes. A possible explanation for this is that the model does not manage to separate the different classes, and therefore uses the most frequent classes as the default answer to be correct most often. See figure 8.

Figure 8: Confusion matrix for the multilabel model. Here it becomes clear how the model favours two of the classes and outright ignores other.



The resulting map made from the model predictions is shown in figure 9. One of the 12 classes is picked to show the random predictions made using this model. The other classes displayed similar random patterns but are not shown here. There is no pattern or grouping of the MiS habitat. A pattern would have been expected since each image tile is only 250 m2.

Figure 9: MiS probability map made from the model in objective one. This map shows the predicted MiS tiles in the test area, and the outline of the reference MiS habitats. Ideally the predictions should only fall inside these lines. Notice the random pattern of predicted MiS habitats.



6.2 Evaluating objective two

The model was evaluated on the test dataset. Classification score during testing of 61%, combined with the result shown in the confusion matrix, made this the best performing model in this thesis. During training this model achieved an accuracy of 78%, but this result was probably caused by overfitting. This performance increase may be attributed to using a pretrained model already knowing some general features, which it found in the tiles. Using only a binary classification method may also have increased the model's performance. The classification report lists the precision, recall, and f1-score of the two classes the model uses. See table 4. The low recall and precision of the MiS class shows that the model is not very good at finding MiS tiles. See figure 10. The f1-score shows that the classifier is not good at finding the MiS tiles, only having an f1-score of 0.15.

Table 4: Evaluation metrics Classification report:				
	precision	recall	f1-score	support
MiS	0.08	0.79	0.15	1121
Not MiS	0.98	0.60	0.74	24104

Figure 10: The precision-recall curve to the model. This curve looks much better than the curve in objective one. The curve tapers off slower and therefore shows that the model is more usable. On the other hand, notice how low the recall has to be to obtain a precision above 80%. This could mean that only obvious image tiles are classified reliably.



Cohen's kappa does give a result of 0.07, like random guessing, not good. The confusion matrix shows that 85 % of the MiS tiles are classified correctly. This is good but we also see that almost half of the non-MiS tiles are classified

as MiS. This makes the model very inaccurate because the total amount of MiS tiles is five times less than what the model finds! See figure 11.

Figure 11: Confusion matrix of the model in objective two showing both raw and normalized data.



The resulting map made from the model predictions is shown in figure 12. Big parts of the MiS habitats are predicted correctly as seen in the map, although large parts of the map are also classified as MiS habitats. This means that the value of this map to a field worker is low since it does only to a small degree narrow down the search space. Figure 12: MiS probability map made from the model in objective two. Here groups and patters of MiS-habitats form. This map shows the predicted MiS tiles in the test area, and the outline of the reference MiS habitats. Ideally the predictions should only fall inside these lines. This could be interpreted as the model is not randomly guessing, like in objective one, but have a understanding for what it is looking for.



6.3 Evaluating objective three

The third objective was to restrict the dataset to only include tiles containing forest. This restriction resulted in a test accuracy of 61%, similar to the model in

objective two. During training, this model achieved an accuracy of 75%. Looking at the confusion matrix shows that the amount of correctly classified MiS tiles it finds have decreased compared to the previous model. Therefore removing the 'noise' from the training data does not create a better model.

The classification report lists the precision, recall, and f1-score of the two classes the model uses. See table 5. The low recall and precision of the MiS class shows that the model is not very good at finding MiS tiles. See figure 13. The f1-score shows that the classifier is not good at finding the MiS tiles, only having an f1-score of 0.12. Lower than the model from objective two.

Table 5: Evaluation metrics objective three Classification report:

	precision	recall	f1-score	support
MiS	0.07	0.60	0.12	1121
Not MiS	0.97	0.61	0.75	24104

Figure 13: The precision-recall curve to the forest mask model in objective three. This curve shows that although the accuracy score of the models in objective two and three are close the model in objective three is worse. Here the area under the curve is lower and the average precision only is 0.6 compared to 0.8.



Cohen's kappa does give a result of 0.05, like random guessing, not good. The confusion matrix shows that 60 % of the MiS tiles are classified correctly. This is worse than the model in objective two. Here are also almost half of the non-MiS tiles classified as MiS. This makes the model very inaccurate because the total amount of MiS tiles is much less. See figure 14.



Figure 14: Confusion matrix with both raw and normalized data for objective three.

The resulting map made from the model predictions is shown in figure 15. Big parts of the MiS habitats are predicted correctly as seen in the map, but large parts are also predicted wrongly. Figure 15: MiS probability map made from the model in objective three. This map shows the predicted MiS tiles in the test area, and the outline of the reference MiS habitats. Ideally the predictions should only fall inside these lines. It seems the model wrongly classifies a bigger area as MiS-habitats compared to the model in objective two.



6.4 Comparison of the objectives

Comparing the three different models one can conclude that the models are bad at predicting MiS tiles. Each objective changed the data that were given to the models, and although it did change the outcome slightly, the quality of the predictions were bad. See the comparison in table 6. Not one of the MiS probability maps can be said to give any aid to a potential user. This thesis goal of creating a usable MiS probability map failed.

Table 6: Comparison of the different models used in each objective. It shows how the model made in objective two performs the best.

	Objective one	Objective two	Objective three
Accuracy score	0.06	0.61	0.61
Cohens kappa	0	0.077	0.045
Precision MiS tiles	-	0.08	0.07
Recall MiS tiles	-	0.80	0.60

7 Discussion

The best model found in this thesis could not reliably be used to find MiS habitats in the test area. The idea to use a CNN's feature extracting ability on forest imagery was rooted in previous results this kind of network had on other image classification tasks. For example, the tree classification done by Trier et al., and the models trained on ImageNet where hundreds of different classes could be classified correctly. Using tiles was thought to be better than only using information from single pixels. This is because the forest may have a structure containing information only found when looking at multiple pixels and their surroundings. An example of structure is a gap in the forest cover which could indicate a fallen tree. Using only single pixel-based models the pixels creating the hole would probably be discarded as noise. A CNN looking at patterns could use the pixels creating the hole as an indicator that it may be an area with fallen trees. Another hypothetical example is finding lush ground vegetation in a forest. Here the trees may have a greater distance between them, compared to other areas, letting more light down to the ground vegetation. The bigger distance between the trees is a structure a CNN could have found. A pixel-based classifier could not have utilized this pattern, potentially losing important information. It would have been interesting comparing the results from a pixel-based classifier with the results from this thesis.

First off all machine learning models need lots of data to produce good results. The hundreds of tiles given to the model during training may have been too low to find a difference between MiS and non-MiS tiles. Here the number of tiles could be increased by producing a bigger dataset, which is expensive, or the tile size could be reduced. Halving the tile size from 50x50 pixels to 25x25 pixels produces four times the tiles to feed the model, but this needs a new model and was not tried out. Secondly, the labels used and regarded as ground truth is based on imperfect field work. Field workers could have missed MiS habitats, classified wrongly or skipped parts of the forest. This is mentioned in the introduction and could potentially confuse the model by giving it wrongly labeled tiles. Finally, and arguably the most important factor in this model's performance, the domain knowledge plays a major role.

7.1 Problems encountered

When working on this project multiple problems were encountered. One of the first problems arose from the size of the image data. Originally the plan was to use the bands from both the VNIR and the SWIR sensors. The raw data was delivered as several stripes of image data from both sensors. Trying to merge the stripes and sensors together into one turned out to be a time-consuming task. Merging one of the thirteen image stripes took an ordinary desktop computer over 30 hours to complete. The time consumption was probably caused by the big file. Therefore, only the VNIR data, which was thought to contain the most useful information, was used. Had more time or a more powerful computer been available the SWIR imagery would be included. Information from this sensor could have provided a better data foundation for the model.

Making the model accept image tiles were also a source of problems. Since the early models were trained on hyperspectral multi-band images, helper functions provided by Keras could not be used. The helper functions enabled easy importing, data argumentation, splitting into training and validation, etc, but does not work with images having more than four bands. This meant that these functions had to be custom made. Time that could have been used creating a better model. When the final models were made using three band tiles, from the PCA dimensionality reduction, the Keras helper functions where used.

Continuing this work I would have investigated the laser data provided in this project. Laser would give the model a different kind of information than hyperspectral data. Knowing the terrain and vegetation structure of an area could maybe give a new model a better chance of finding MiS habitats. Using both laser- and hyperspectral data as input into a model is one possibility. I would also have tried different tiles sizes. Maybe a bigger tile could contain more information and patterns the model could have used. Or maybe it is the other way that a small tile could have made the differences between habitats more pronounced. Lastly, it would have been interesting to explore more what kind of tiles the model managed to predict right. Do these tiles have something in common?

8 Conclusion

The goal of this thesis was to explore the possibilities of using machine learning in the form of a CNN, to find MiS habitats in hyperspectral images. Unfortunately, the models created failed to produce meaningful results. During training, the best performing model had an accuracy score of 61%, but when evaluated on the test data it became clear its usability was bad. The confusion matrix showed that 85% of the MiS tiles are classified correctly. This is good but we also see that almost half of the non-MiS tiles are classified as MiS. This makes the model very inaccurate because the total amount of MiS tiles is five times less than what the model finds. This means that the value of this map to a field worker is low since it does only to a small degree narrow down the search space. The model did classify many of the MiS habitats correctly, but the problem was that it classifies a lot more wrong. Giving the CNN more and varied data could give better results. Better knowledge of the data could have improved the model performance in several ways. Knowing what features characterises the different MiS habitats may have helped tailor the model to be more suitable for this specific problem. Another advantage is knowing what parts of the data is valuable, and how to separate it from noise. Trying a different approach would also change the model performance. Instead of using the image tiles, using bigger or smaller part of the aerial photographs would effect how the model behaves. The conclusion of this thesis is that the method used here did not work, but more domain knowledge and a different approach may give different and better results.

9 Acknowledgments

I wish to thank several people for their contribution to this thesis. Odd Arne Brenn and Gyrd Harstad from Mjøsen Skog, for teaching and showing me how practical fieldwork is done. This gave me insight into what signs experienced woodsmen look for when surveying the forest. Special thanks to Dr. Hans Ole Ørka, my thesis supervisor for his guidance and valuable support. Meeting him weekly gave me frequent and valuable input on problems and his thoughts helped shape this thesis for the better. Tusen takk!

10 Bibliography

- [1] Landbruksdirektoratet. Miljøregistreringer i skog Landbruksdirektoratet, Feb 2019. [Online; accessed 12. Feb. 2019]. URL: https://www.landbruksdirektoratet.no/no/eiendom-og-skog/ skog-og-miljoregistreringer/miljoregistreringer.
- [2] Landbruks og matdepartementet. The Norwegian Forest Habitat Inventory. Regjeringen.no, Jun 2010. URL: https://www.regjeringen.no/no/ dokumenter/The-Norwegian-Forest-Habitat-Inventory/id610571.
- [3] Eva L. Eriksen, Heidrun A. Ullerud, Rune Halvorsen, Sigrun Aune, Harald Bratli, Peter Horvath, Inger K. Volden, Anders K. Wollan, and Anders Bryn. Point of view: error estimation in field assignment of land-cover types. *Phytocoenologia*, pages -, 11 2018. URL: http://dx.doi.org/10.1127/ phyto/2018/0293, doi:10.1127/phyto/2018/0293.
- [4] Yunsuk Kim, Zhiqiang Yang, Warren B. Cohen, Dirk Pflugmacher, Chris L. Lauver, and John L. Vankat. Distinguishing between live and dead standing tree biomass on the north rim of grand canyon national park, usa using small-footprint lidar data. *Remote Sensing of Envi*ronment, 113(11):2499 – 2510, 2009. URL: http://www.sciencedirect. com/science/article/pii/S0034425709002338, doi:https://doi.org/ 10.1016/j.rse.2009.07.010.
- [5] Raita Säynäjoki, Petteri Packalen, Matti Maltamo, Mikko Vehmas, and Kalle Eerikainen. Detection of aspens using high resolution aerial laser scanning data and digital aerial images. *Sensors (Basel)*, 8(8):5037-5054, Aug 2008. URL: https://www.ncbi.nlm.nih.gov/pubmed/27873799, doi:10.3390/s8085037.
- [6] Raymond F. Kokaly, Barnaby W. Rockwell, Sandra L. Haire, and Trude V.V. King. Characterization of post-fire surface cover, soils, and burn severity at the cerro grande fire, new mexico, using hyperspectral and multispectral remote sensing. *Remote Sensing of Envi*ronment, 106(3):305 - 325, 2007. URL: http://www.sciencedirect. com/science/article/pii/S0034425706003026, doi:https://doi.org/ 10.1016/j.rse.2006.08.006.
- [7] Matti Maltamo, Erik Næsset, and Jari Vauhkonen. Forestry Applications of Airborne Laser Scanning: Concepts and Case Studies. Springer, 01 2014. doi:10.1007/978-94-017-8663-8.
- [8] Forest: deadwood, May 2019. [Online; accessed 2. May 2019]. URL: https: //www.eea.europa.eu/data-and-maps/indicators/forest-deadwood/ forest-deadwood-assessment-published-may-2010.
- [9] Laurel Ballanti, Leonhard Blesius, Ellen Hines, and Bill Kruse. Tree species classification using hyperspectral imagery: A comparison of two classifiers. *Remote Sensing*, 8:445, 05 2016. doi:10.3390/rs8060445.
- [10] Kar-Han Tan and Boon Pang Lim. The artificial intelligence renaissance: Deep learning and the road to human-level machine intelligence. APSIPA

Transactions on Signal and Information Processing, 7, 07 2018. doi:10.1017/ATSIP.2018.6.

- [11] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. Cnn features off-the-shelf: An astounding baseline for recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2014. [Online; accessed 8. May 2019]. URL: https://www.cv-foundation.org/openaccess/content_cvpr_workshops_2014/W15/html/Razavian_CNN_Features_Off-the-Shelf_2014_CVPR_paper.html.
- [12] ImageNet Wikipedia, Mar 2019. [Online; accessed 27. Mar. 2019]. URL: https://en.wikipedia.org/wiki/ImageNet.
- [13] Øivind Due Trier, Arnt-Børre Salberg, Martin Kermit, Øystein Rudjord, Terje Gobakken, Erik Næsset, and Dagrun Aarsten. Tree species classification in norway from airborne hyperspectral and airborne laser scanning data. *European Journal of Remote Sensing*, 51(1):336-351, 2018. URL: https://doi.org/10.1080/ 22797254.2018.1434424, arXiv:https://doi.org/10.1080/22797254. 2018.1434424, doi:10.1080/22797254.2018.1434424.
- [14] Steven E. Franklin and Oumer S. Ahmed. Deciduous tree species classification using object-based analysis and machine learning with unmanned aerial vehicle multispectral data. *International Journal of Remote Sensing*, 39(15-16):5236-5245, 2018. URL: https://doi.org/10.1080/01431161.2017.1363442, arXiv:https://doi.org/10.1080/01431161.2017.1363442, doi:10.1080/01431161.2017.1363442.
- [15] R. H. J. Heim, I. J. Wright, H.-C. Chang, A. J. Carnegie, G. S. Pegg, E. K. Lancaster, D. S. Falster, and J. Oldeland. Detecting myrtle rust (austropuccinia psidii) on lemon myrtle trees using spectral signatures and machine learning. *Plant Pathology*, 67(5):1114–1121, 2018. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/ppa.12830, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/ppa.12830, doi:10.1111/ppa.12830.
- [16] Kai Wang, Steven E. Franklin, Xulin Guo, and Marc Cattet. Remote sensing of ecology, biodiversity and conservation: A review from the perspective of remote sensing specialists. *Sensors*, 10(11):9647–9667, 2010. URL: http: //www.mdpi.com/1424-8220/10/11/9647, doi:10.3390/s101109647.
- [17] Figure 3. Example of a hyperspectral image data cube. Each pixels consists of a complete reflection spectrum at its position., Aug 2018. [Online; accessed 15. Feb. 2019]. URL: https://www.researchgate.net/figure/ Example-of-a-hyperspectral-image-data-cube-Each-pixels-consists-of-a-complete-reflect fig2_259215766.
- [18] Jürgen Schmidhuber. Deep learning in neural networks: An overview. Neural Networks, 61:85 - 117, 2015. URL: http://www.sciencedirect. com/science/article/pii/S0893608014002135, doi:https://doi.org/ 10.1016/j.neunet.2014.09.003.

- [19] 7wf.jpg (JPEG Image, 1320 by 708 pixels), Oct 2018. [Online; accessed 20. Feb. 2019]. URL: https://i.vas3k.ru/7wf.jpg.
- [20] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. CoRR, abs/1311.2901, 2013. URL: http://arxiv.org/ abs/1311.2901, arXiv:1311.2901.
- [21] Transfer learning Wikipedia, Mar 2019. [Online; accessed 27. Mar. 2019]. URL: https://en.wikipedia.org/wiki/Transfer_learning.
- [22] Pedro Marcelino. Transfer learning from pre-trained models. Towards Data Science, Oct 2018. URL: https://towardsdatascience.com/ transfer-learning-from-pre-trained-models-f2393f124751.
- [23] Specifications, Feb 2019. [Online; accessed 28. Feb. 2019]. URL: https: //www.hyspex.no/products/all_specs.php.
- [24] Ground sample distance Wikipedia, Mar 2019. [Online; accessed 7. Mar. 2019]. URL: https://en.wikipedia.org/wiki/Ground_sample_ distance.
- [25] Keras Wikipedia, Feb 2019. [Online; accessed 28. Feb. 2019]. URL: https://en.wikipedia.org/wiki/Keras.
- [26] GDAL/OGR contributors. GDAL/OGR Geospatial Data Abstraction software Library. Open Source Geospatial Foundation, 2018. URL: http://gdal.org.
- [27] Principal component analysis Wikipedia, Mar 2019. [Online; accessed 27. Mar. 2019]. URL: https://en.wikipedia.org/wiki/Principal_ component_analysis.
- [28] QGIS Wikipedia, Feb 2019. [Online; accessed 28. Feb. 2019]. URL: https://en.wikipedia.org/wiki/QGIS.
- [29] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825– 2830, 2011. URL: https://scikit-learn.org/stable/auto_examples/ model_selection/plot_precision_recall.html.



Norges miljø- og biovitenskapelige universitet Noregs miljø- og biovitskapelege universitet Norwegian University of Life Sciences

Postboks 5003 NO-1432 Ås Norway