Aalborg Universitet



Machine learning for automatic detection of historic stone walls using LiDAR data

Trotter, Ezra Francis Leslie; Fernandes, Ana Cristina Mosebo; Fibæk, Casper Samsø; Keßler, Carsten

Published in: International Journal of Remote Sensing

DOI (link to publication from Publisher): 10.1080/01431161.2022.2057206

Creative Commons License CC BY 4.0

Publication date: 2022

Document Version Publisher's PDF, also known as Version of record

Link to publication from Aalborg University

Citation for published version (APA): Trotter, E. F. L., Fernandes, A. C. M., Fibæk, C. S., & Keßler, C. (2022). Machine learning for automatic detection of historic stone walls using LiDAR data. International Journal of Remote Sensing, 43(6), 2185-2211. https://doi.org/10.1080/01431161.2022.2057206

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.





International Journal of Remote Sensing

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/tres20

Machine learning for automatic detection of historic stone walls using LiDAR data

Ezra Francis Leslie Trotter, Ana Cristina Mosebo Fernandes, Casper Samsø Fibæk & Carsten Keßler

To cite this article: Ezra Francis Leslie Trotter, Ana Cristina Mosebo Fernandes, Casper Samsø Fibæk & Carsten Keßler (2022) Machine learning for automatic detection of historic stone walls using LiDAR data, International Journal of Remote Sensing, 43:6, 2185-2211, DOI: 10.1080/01431161.2022.2057206

To link to this article: https://doi.org/10.1080/01431161.2022.2057206

© 2022 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



6

Published online: 08 Apr 2022.

Submit your article to this journal 🗗





View related articles 🖸



則 🛛 View Crossmark data 🗹



👌 OPEN ACCESS 🚺

Check for updates

Machine learning for automatic detection of historic stone walls using LiDAR data

Ezra Francis Leslie Trotter^{a,b}, Ana Cristina Mosebo Fernandes (D^{a,b}, Casper Samsø Fibæk^{a,b} and Carsten Keßler^{a,c}

^aDepartment of Planning, Aalborg University, Copenhagen, SV, Denmark; ^bNIRAS Consulting, Sortemosevej, Allerød, Denmark; ^cDepartment of Geodesy, Bochum University of Applied Sciences, Bochum, Germany

ABSTRACT

Stone walls in the landscape of Denmark are protected not only for their cultural and historical significance but also for their vital role in supporting local biodiversity. Many stone wall structures have either disappeared, suffered substantial damage, or had segments removed. Additionally, as it stands today, the registry of these structures, managed by each municipality, is outdated and incomplete. Leveraging recent developments in Machine Learning and Convolutional Neural Networks (CNNs), we analyze the publicly available terrain data (40 cm resolution) derived from the Danish LiDAR data, using a U-Net-like CNN model to assess the stone walls dataset and provide for an update of the registry. While the Digital Terrain Model (DTM) alone provided good results, better results were obtained when adding Height Above Terrain (HAT) and an additional DTM layer with a Sobel filter applied. Using a pixel-wise evaluation, there was an overall agreement of 93% between ground truth and prediction of stone walls in a validation area and 88% overall agreement for the whole predicted area. Good generalizability was found when externally validating the model on new data, showing positive results for both the existing stone walls and predicting new potential ones upon visualisation. The method performed best in open areas, however positive results were also seen in forested areas, although denser areas and urban areas presented as challenging. Given the lack of a reference dataset or other studies on this specific matter, the evaluation of our study was heavily based on the stone walls registry itself complemented by visual inspection of the predictions and on the ground in the Danish municipality of Ærø. Automating the process of identifying and updating the stone walls registry in Denmark is of great relevance to the local governments. We suggest the development of a Decision Support System to allow municipalities access to the results of this method.

ARTICLE HISTORY

Received 20 September 2021 Accepted 18 March 2022

KEYWORDS

Stone walls; LiDAR; DEM; deep learning; archaeology

1. Introduction

Stone walls are structures in the Danish landscape that are protected for their cultural and historical significance. The oldest structures were built in the first century, most commonly as property boundaries and to mark administrative divisions. The most recent ones

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/ licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

CONTACT Ana Cristina Mosebo Fernandes 🔯 afer@niras.dk 🖃 Department of Planning, Aalborg University, A.C. Meyers Vænge 15, Copenhagen, SV 2450, Denmark

^{© 2022} The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

2186 😉 E. F. L. TROTTER ET AL.

were built in the 1800s and were used to divide and mark agriculture fields and to delineate forests and woods planted by royal decree (Kulturministeriet 2009). In addition to their intended purpose, stone walls also serve as vital green corridors between land-scapes (Pedersen 2021). Therefore, these structures are important to preserve for their historical importance and their vital role for local biodiversity.

As a consequence of urban expansion and development of agricultural land, a substantial number of Denmark's stone wall structures have either already disappeared or have suffered considerable damage (Kulturministeriet 2009). Danish stone walls and sand walls have therefore been classified as protected in 1992 by the nature protection law (Kulturministeriet 2004), which was followed by an update in 2004 with the "Museums' law", which transferred responsibility to the individual municipalities. Today, each municipality is responsible for managing and protecting its culturally significant stone walls and sand walls. However, as it stands today, the Danish Ministry of Culture's aggregate dataset is not up to date and does not fully account for all the stone walls and sand walls in each municipality (Christensen 2020), triggering a need to research methods to update the registry in an automated fashion.

The definition of stone walls by the law characterises them as 'Man-made, linear elevations of stone, earth, turf, seaweed or similar materials which function or have functioned as fences and have or have had the purpose of marking administrative property or use boundaries in the landscape' (Kulturministeriet 2009). Protected stone walls include the structures falling under this definition and those already registered in the 1:25 000 topographic maps (SDFE. 1977–1992), in the public domain and those situated on or near protected habitats. Their physical characteristics vary in size, shape and materials. Generally, the walls are between 0.5 and 1.5 m in height, 1.5 m in width; made of either stone, heather peat, soil or a combination thereof (Figure 1).



Figure 1 Different shapes and forms of stone walls and sand walls, in its original form (a and b) and how they look like today (Kulturministeriet 2009).

The removal or partial destruction of protected stone walls is against the established law (Kulturministeriet 2004), however, it is possible to apply for a dispensation to the respective municipality, who will assess the case for removal or alteration. For this, an upto-date registry is necessary to verify the structure in question and monitor and ensure the preservation of the protected stone walls.

With this in mind, this study aims to analyze terrain data derived from the Danish LiDAR (Light Detection and Ranging) dataset to assess the stone walls dataset and provide for an update of the registry. For updating the registry, the study will focus on two main tasks: (1) Analyze stone walls with the terrain data, and identify stone walls or segments of walls that no longer are existent on the ground and were removed; (2) Find and map potential new stone wall structures that are not registered but should be included in the registry.

To accomplish this, we will analyze LiDAR derived data in order to profile the stone walls in Ærø, then we will use a simplified morphometric algorithm to identify their topographic peak points. This analysis will also validate the stone walls dataset and serve as feature-engineering to prepare the dataset for the second task. We will apply a Deep Learning (DL) method by using a Convolutional Neural Network (CNN) model in order to find and map potential stone walls that are not registered, using the validated stone walls dataset. With this, we intend to provide an automated method for updating the registry of protected stone walls, which coupled with a visualisation tool and onsite confirmation by experts, will provide the support necessary to the Danish municipalities.

The rest of the study is structured as followed: Section 2 references previous work related to the data and methods used in similar contexts to this study; Section 3 describes the materials and methods utilised for accomplishing this work, namely the validation of the stone walls dataset and its pre-processing, the U-Net-like model for the prediction of new stone walls, as well as its post-processing; Section 4 provides the results obtained and their assessment; Section 5 engages in the discussion of the results and its applications, and finally Section 6 summarises the work performed in this study.

2. Related work

2.1. LiDAR data

While stone walls are prominent features of the Danish landscape, it is difficult to observe them when they are located in woods and forests. Therefore, elevation models derived from LiDAR data provide a valuable means to identify structures and objects on the surface since they are not affected in the same way by vegetation due to the ability of LiDAR to penetrate forest and scrub canopies (Chen, Gao and Devereux 2017). A myriad of Archeology mapping studies uses derived elevation models from LiDAR data to identify objects and structures on the topographic landscape (Øivind, Cowley and Waldeland 2019; Chase et al. 2012; Guyot et al. 2021).

2.2. Deep learning

Fully Convolutional Neural Networks (FCNs) were originally developed for semantic segmentation of medical images. U-net, a popular architecture, was first introduced in 2015 (Ronneberger, Fischer and Brox 2015). Maxwell et al. (2020) have used semantic

segmentation to detect surface disturbance caused by mining from topographic maps, where the study mentions a high level of accuracy achieved. The study uses a modified U-Net architecture, which classifies each pixel location separately. The output is the probability of each pixel belonging to each of the classes defined in the model. This method differs from instance segmentation, where the goal is not only to identify the object from the background, but to identify and attribute a label to each pixel, as well as an individual object of a label. Maxwell, Pourmohammadi and Poyner (2020), proposes instance segmentation to map topographic features (valley fill faces) using LiDAR-derived data. The study explores the application of Mask *R*-CNN and mentions successful accuracies with such method.

3. Data and methodology

Data quality is essential in achieving good model performance and significant results. While dealing with datasets representing real-world applications, it is all too common to have noisy and faulty data, often necessitating rigorous pre-processing. Neglecting the importance of data processing and preparation can lead to data cascades (Sambasivan et al. 2021), where data issues cause downstream effects, leading to poor model performance and output. In our study, the stone walls dataset represents the stone wall structures present in the Danish landscape. The guality of the raw dataset is questionable; it is neither an accurate representation the stone walls' precise geographic locations nor is the dataset current with the ground truth. These dataset characteristics determine the design and structure of the project and the approach to the problem statement. In order to update the stone walls registry, we will first engage in verifying the presence of the individual walls against the Danish elevation model, since most of the structures are observable in the terrain data. Here we intend to remove sections of stone walls that no longer exist, and adjust the position of the walls in the dataset in order to reflect their actual location on the ground better. We will then use the validated data to detect and map potential non-registered stone walls, by performing a regression task using a Deep Learning model with a U-Net-like architecture.

3.1. Terrain data

A digital terrain model (DTM) (SDFE. 2014) and a digital surface model (DSM) (SDFE. 2014) were used in this study, both of which were downloaded from The Danish Map supply's website and are based on aerial laser scanning measurements taken in December of 2014. Both datasets describe Denmark's surface in relation to mean sea-level and are provided as raster layers at 0.4 m resolution, covering the entirety of Denmark. Each individual pixel value is precise to 15 cm horizontally and 5 cm vertically, and were made available in the ETRS89 UTM 32N coordinate system (EPSG: 25,832) (SDFE. 2020a).

A third dataset was created by subtracting the pixel-wise value of the DSM from that of the DTM to give the Height Above Terrain (HAT), also known as normalised DSM (nDSM) (Chen, Gao and Devereux 2017). This was done over the extent of each raster, thereby creating a new raster of the same size, with pixel values 'HAT = DSM – DTM'. The HAT can

describe the height of structures above the surface, providing height information on individual objects, such as buildings, and in this case on stone wall structures. The analysis of the HAT alone was not sufficient to identify stone walls, as they proved hard to distinguish when covered by vegetation or dense forest. However, given its potential of providing additional context of the stone wall structures' location, it was considered as an additional layer to the training data.

Additionally, a Sobel filter was used in conjunction with the DTM to create a new layer. The Sobel operation implements a 2D spatial gradient calculation on an image by sliding a pair of convolution masks (3x3) on the x-direction (horizontal) and on the y-direction (vertical), respectively (Vincent and Folorunso 2009). The mask manipulates the pixels one by one, changing the value of the pixel according to the kernels. The Sobel layer was created in Python using the Buteo toolbox, accessible through the project repository https://github.com/casperfibaek/buteo. Such layer was added given its usefulness in edge detection to improve the identification of stone walls by the model.

3.2. Study site and stone wall dataset

A smaller study area was selected to reduce the amount of data required and overall processing time. The island Municipality of Ærø lies in the Baltic Sea between the Danish Island of Funen and the German Region of Schleswig and has an area of 88 km² (Figure 2). It was chosen for its smaller size relative to other municipalities, its abundance of stone walls, and its gently undulating landscape.



Figure 2 The study site: the municipality of ærø.

2190 😉 E. F. L. TROTTER ET AL.

A digitised map of Denmark's protected stone walls is made available as part of the Danish Ministry of Culture's stone wall registry. This map was first digitised on the 1st of July 1992 and updated in 2006, and is made publicly available as a vector dataset for download through the Ministry's data portal (Kulturstyrelsen, Slots- og. 2011). On Ærø, the dataset contains 2.766 stone walls for a total length of around 514 kilometers (Figure 3). Each wall in the dataset is represented as a vector linestring and the walls' associated metadata. The metadata contains information such as the date of registry, the walls' current condition and the institution responsible for the data integrity.

The majority of stone walls on Ærø were registered as a result of the digitisation of the 1:25 000 topographic map of Denmark. Most of the metadata is either incomplete or no longer current. This is especially critical in some areas where almost 30% of the protected stone walls did not appear in the subsequent versions of the topographic map (Kulturstyrelsen, Kultur Ministeriet – Slots og. 2020).

An analysis of the 2012 Corine Land Cover (CLC) (SDFE. 2012) sourced from SDFE's data portal reveals that stone walls on Ærø are found in generally similar types of landcover as those in the rest of Denmark. This analysis found that ~80% of the stone walls in the data set were located on agricultural land, both on Ærø and in the rest of Denmark. This similarity continues with ~5% of stone walls being located in discontinuous urban fabric, both on Ærø and Denmark (Figure 4). Where Ærø differs is in the number of stone walls found in forested areas. The municipality has little forested land, and as a result, almost no stone walls are found in forests, whereas ~15% of Denmark's stone walls are found in forested areas. For this analysis, all landcover types where a stone wall is present were included.



Figure 3 Stone walls and DTM on Ærø.



Figure 4 Land cover types where stone walls are located. Denmark (top), Ærø (bottom).

3.3. Pre-Processing of the stone wall dataset

The initial step was to validate the stone wall dataset against the most recent DTM data. The Digital Terrain Model is useful for detecting stone walls because of their topographic characteristics, standing out from their surrounding landscape in a hillshade analysis (Figure 5).

This initial step is deemed necessary given the difference between the two datasets regarding their production date. The DTM data was current as of 2014, while the stone wall reference dataset was digitised in 2006. As discussed earlier, when the reference was digitised there was little effort to validate the individual walls. Reference data representation accuracy has a large effect on the overall performance and results of a DL model (Goodfellow, Bengio and Courville 2016), and it was important to find and remove stone walls segments that had been removed or altered in the intervening time to achieve the best possible results during the later stages of our study.



Figure 5 DTM in hillshade with stone walls dataset shown in red.

The initial inspection and validation had four distinct steps:

- (1) Create profiles along each line representing each stone wall
- (2) Check each profile for the presence of a stone wall
- (3) Redraw the dataset with absent walls removed
- (4) Validation of the corrected dataset

Step 1: Creating Profiles

The initial step taken was to segment each line in the dataset into 5 m sections to create a profile at the end of each section. Some walls were represented as straight linestrings, and others were multi-linestring objects. The presence of multi-linestring objects required an initial step of segmenting each multi-linestring into its composite lines. This was necessitated by the later process of recreating the dataset in the adjusted positions. At each 5 m subsection, a tangential line is created (cross-section) with a length of 10 m, which is then broken up into 0.4 m subsections, corresponding to the pixel resolution of the DTM. At each of these points, the elevation value of the DTM is extracted, and a 3D multipoint object was created containing the x and y positional coordinates of the points in each profile, along with the value for the elevation (Figure 6). The initial dataset of 2.766 lines yielded 113.089 profiles, each comprising of 50 points.

Step 2: Identifying Stone walls

The datasheets provided by the Culture Ministry provide a reference of which dimensions a stone wall can have. Additionally, by plotting and completing a visual inspection, it is observable whether or not a profile represents a wall. With a small number of walls or profiles, it would be possible to sort the profiles manually. However, due to the large number of profiles created in Ærø alone, an automated method was necessary. The 'find peaks' function of the SciPy Python package (Virtanen, Gommers and Oliphant et al. 2020), originally intended for use in Computational Biology and Bioinformatics, was utilised to identify 'peaks', which in our case were the peaks of the walls. This method was extremely successful at identifying stone walls that had a prominent peak, which was the case for



Figure 6 Examples of stone wall profiles. The top three profiles show the presence of stone walls, while the bottom three are examples of profiles where a wall is no longer present.

the bulk of the dataset. However, walls that were damaged, partially removed or covered with earth over time were more challenging for the algorithm to correctly categorise. Walls of such type were marked with an 'unclear' category for later manual inspection, while the remaining walls were designated either the 'wall' or 'not wall' category.

Additionally, the 'find peaks' function stored the index of each peak. This allowed for the calculation of the position of the peak in relation to the cross-section. If multiple peaks were found, only the peak that was closest to the center of the profile was stored. These peak indexes were used to adjust the position of the linestrings in the stone walls dataset to the stone walls actual position on the ground in the following step.

Step 3: Rebuilding the Dataset

Firstly, it was necessary to recreate the dataset with only the sections that were categorised as a wall. This was achieved by redrawing the linestrings after their categorisation. For an individual segment of stone wall, starting at the first profile, a linestring was drawn between the peak of the current wall and the peak of the following wall. This was performed only in the case that the current and next wall were categorised as walls. In this way, the entire dataset was redrawn with the non-wall sections removed and with the remaining walls adjusted to their actual position as reflected in the DTM. This provided for the feature engineering of the dataset, in order to maximise the extraction of features suited to represent the target data for the CNN model training (He, Zhao and Chu 2021).

Step 4: Validation

For the majority (>90%) of the profiles, the presence of a stone wall was easily identified. However, for the cases where the wall was either very low, or partially removed, the profiles could not be reliably classified by our simple algorithm. As described above, for any profile not quickly identified (the prominence of the peak was less than 0.3 m), a third class 'unclear' was attributed. In lieu of developing a more robust and complex

algorithm, it was necessary to manually inspect these profiles with a visualisation software tool (QGIS 3.16 (QGIS.org 2021)). This inspection was done by visually comparing the 'unclear' profiles against the DTM in hillshade, coupled with the profile tool plugin. Any profiles that suggested the wall no longer was present or if verification was not possible using these methods, were removed from the dataset.

3.4. Data preparation and image patches generation

The next step was to train a CNN on the updated stone wall dataset. For this, the stone wall would first need to be converted into a raster format. The DEM data, which would be used as training data, had a spatial resolution of 0.4 m per pixel; therefore the stone walls dataset would need to be rasterised to match this resolution. During the transformation to raster format, a down-sample of pixels was first performed (down-sampled to 10 cm) to expand the number of presence pixels. The pixels were then restored back to the target resolution of 0.4 m, resulting in a 'anti-aliased' walls dataset (example on Figure 7). All the pixels covering the location of the wall were given a float value ($0 < x \le 1$ 1), commensurate to the distance from the center of the pixel to the center of the wall. Pixels that intersected the centreline of the wall were given a value 1.0. This transformation helped to account for spatial uncertainty but also gave flexibility to the output prediction. In doing this, we turn the problem into one of regression to compensate for the number of absence pixels present in the patches while extracting the target data for training the model. A high number of absence pixels can occur due to the relationship between the buffer distance and the size of the patches (64x64), where some patches contained no stone wall segments.

Deep Learning models that employ CNNs require the training data to be in the form of small patches since the spatial context information is learned by filters. Additionally, labels must be provided along with associated environmental data. In our case, the rasterised stone walls acted as the labels, and the stacked DTM, HAT and Sobel filter were the associated image data. Using scripts produced in Python, the training data was broken into 64×64 pixel patches, resulting in 24.782 patches of stacked DEM data and the same amount of associated rasterised stone wall labels.



Figure 7 (a) example of an "anti-aliased" stone wall; (b) example of the location and size of the 64×64 pixel patches used for the label data.

Given the size of the training data, and to increase generalizability and decrease overfitting within the model, Data Augmentation was used (Shorten and Khoshgoftaar 2019). Data Augmentation is the process of increasing the amount of available data by adding manipulations to the 'real' dataset to create additional data that is similar – this new data is called 'synthetically modified data'. Specifically, we implemented a geometric transformation by rotation augmentation on each patch, such that each patch was rotated 0, 90, 180, and 270 degrees. In the case of our data, this method is considered a 'safe' augmentation transformation, given its like-lihood in preserving the label (Shorten and Khoshgoftaar 2019), because the assumption was that stone wall structures do not have a specific orientation. After splitting the dataset into train and test sets, the Data Augmentation process was applied to the train set only. Such transformations can help reduce overfitting by creating more training data (Shorten and Khoshgoftaar 2019).

For the image inputs, patches were extracted in the buffered area around the rasterised stone wall dataset. Additionally, to investigate the effect that absence data had on our results, additional absence data was added to the final dataset, which was created by extracting patches from other areas on Ærø. Because it was initially theorised that the model would have the most difficulty in urban areas and in differentiating modern walls from historic stone walls, 9.307 patches of absence data were added to the dataset from urban areas. This was done by analyzing the land cover of the island using a CORINE landcover type layer and extracting the patches from urban areas where no protected walls were located. Once the absence data was added, the initial dataset of size 24.782 × 64x64 was increased to size $33.819 \times 64x64$.

When training the model, the dataset was shuffled and then split in train and test sets in a 70–30% ratio, respectively, using Scikit-Learn's function 'train_test_split' (Pedregosa et al. 2012). For the train set, after augmenting the data with rotations, the final size of the dataset was 94.692 × 64x64.

The inherent imbalance in the dataset, where each patch contained many more 0' (no wall) pixels than the number of 1' (wall) pixels, caused difficulties with the loss function. Mean Squared Error (MSE) was used as the model's loss function, and given its calculation method, the loss values ended up being very small. In this way, the input data was scaled to prevent the optimizer to set all the weights to zero.

3.5. Model training and prediction

The creation and training of the DL model was completed in Python, using the Keras API (Chollet 2015) for TensorFlow (Abadi et al. 2016) to create the model.

For this step, A Fully Convolutional Neural Network (FCCN) model design was used, as this allowed for the output of a prediction raster of the same size as the input (64x64 pixels). The model used has a U-Net-like architecture, with initial down-sampling followed by up-sampling. The U-Net model has an expansive path symmetric to the contracting path, leading to a U-shaped architecture. After each down-sampling convolution (a 3×3 convolution followed by a rectified linear unit (ReLU) and a 2×2 max pooling layer of stride 2), a skip connection (concatenation of a corresponding layer between the contracting and expansive path) is performed to provide information of localisation accuracy, reduced by the use of max-pooling layers (Ronneberger, Fischer and Brox 2015). In this

study, the model architecture followed a similar architecture to that of U-net, employing 6 down-sampling blocks, each containing two convolution layers of 3×3 -sized filters, where each block was followed by a 2×2 max pooling layer of stride 2 and with 0' padding. The three down-sampling blocks used 32, 64 and 96 filters each respectively. Two 3×3 transposed convolution layers were used for the expansive path, with 64 and 96 filters, respectively, followed by a concatenation of the previous transposed layer with a layer from the expansive path (see Figure 8 for a visualisation of the model architecture). The model accepted an input vector of size $64 \times 64x3$, one channel each for the layers DTM, HAT and DTM with Sobel filter. The output was a vector of size $64 \times 64x1$. A ReLU activation function was used in the final layer, where each individual pixel within the patch was designated with a value higher than 0, corresponding to the likelihood that each pixel contained a stone wall. Each pixel corresponds to an area of $0.4 m^2$ on the original raster.

Besides the lower number of blocks and filters used in both the contracting and expansive paths, the activation function used in the convolution layers differed from the original U-Net architecture. Instead of the ReLU function, a similar function named Swish (Ramachandran, Zoph and Le 2018) was applied, after testing the model through various iterations with both, and the latter resulted in better overall model performance. The adaptive moment estimation (Adam) optimisation was used (Kingma and Ba 2015), and a callback was defined, which included a step decay for the learning rate, initialised at 0.001, in order to optimise learning and lead the model to quickly converge to a good solution, and Early Stopping, by monitoring the validation loss function during training. For more details on the specification of the model parameters, we refer to the source code at the associated GitHub repository¹.



Figure 8 Model architecture of the CNN model with a shallower structure compared to the U-Net original model.

The optimisation of hyperparameters is important for achieving the best possible results with a CNN (LeCun, Bengio and Hinton 2015; Goodfellow, Bengio and Courville 2016)). With an FCCN, the most significant being the number of filters used in each convolution layer, the size of the kernel that executes the filter, the size of the 'stride' as the kernel moves across the patch, and the padding for the cases where the kernel area goes off the edge of the patch. A complete grid search of the entire hyperparameter space was not possible due to time constraints, however after narrowing down to a certain number of parameters, the final set of values previously described was defined after iterating the model 5 times, averaging the results and selecting the values with the best overall performance. This evaluation method also includes the additional parameters previously described.

The loss function used was the Mean Squared Error (MSE). The MSE loss function minimises the squared differences between the estimated and the target values and is one of the most typically used for regression problems (Carvalho et al. 2018). According to (Lathuilière et al. 2020) study, the choice of a loss function to use depends on dataset and the model architecture, and it is recommended to try different losses at the initial stage of the process. During training, four other losses were used to run the model, namely the Mean Absolute Error (MAE), Huber loss, log cosh, and MSE. We iterated the model training multiple times over the same loss function and adapted the parameters to understand the impact of the loss on the results.

After model training, a prediction could be made using stacked DTM, HAT and Sobel filter data as inputs, outputting a series of patches representing the predicted areas of stone walls. The output was in the form of a 3-Dimensional array which was then converted to raster format using the Buteo Toolbox in Python. During the prediction, offsets were applied on the predicted patches in order to reduce the noise produced by the patches borders during the extraction of the images. A sequence of offsets was used: 1) 16×16 , 2) 32×32 , and 3) 48×48 ; merged with the median.

3.6. Assessment and post-processing

The lack of studies referencing this stone walls dataset and the dataset's inaccuracy and overall poor condition presents a serious challenge when comparing our results. Because of this, the results of our analysis are compared with the original dataset, where the assumption is made that the locations of the stone walls here are correct.

The assessment of the results was undertaken using a combination of quantitative and qualitative analysis. During a quantitative assessment analysis, the performance and results from the loss function (MSE), additional metrics (MAE and RMSE), and the validation loss were analyzed. This was done after iterating the model multiple times and averaging the results. An assessment based on the pixel value of the prediction was done by comparing the true values and the predicted values for each patch. Taking the intersection of the presence pixels in the input and prediction and the intersection of the absence pixels and dividing by the number of pixels, giving a number between 0 and 1, a representation of the overall accuracy of the prediction was calculated.

In order to obtain a more accurate assessment of the predictions on the true values, a small validation area was selected. This area, as best as could be ascertained given the dataset, reflected the ground truth. The values of the pixels for the selected area 'ground truth' were compared with the predicted values using the same pixel-wise metric. 2198 👄 E. F. L. TROTTER ET AL.

The qualitative analysis was focused on the assessment based on the visualisation of the predicted images. The results from the different training runs were inspected and compared with the original stone walls dataset to verify: 1) the prediction of stone walls in relation to the existing structures; 2) new predicted walls; 3) walls or segments of walls that were no longer present, therefore not predicted; 4) prediction errors and false positives; 5) general noise present in the predictions, surrounding the predicted values and the absence areas. As a final assessment, a field inspection was conducted by first selecting cases of stone walls relevant to verify. In this way, a verification of predicted stone walls and a comparison with the ground truth was made to provide additional validation.

The output of the model shows predicted wall locations in the study area, where pixels with a value >0 indicated the presence of a stone wall. The final step of the analysis is to separate the walls into their respective categories: 1) walls that appear both in the initial dataset and the prediction 2) walls that appear in the initial dataset and not in the prediction, 3) walls that are not present in the initial dataset but do appear in the prediction.

In order to separate the new and removed walls from the dataset, a comparative analysis was undertaken between the initial and prediction datasets. Removing walls that appear in both the prediction and the original wall dataset from the prediction raster reveals the new walls that have been 'discovered' by the algorithm. Removing walls that appear in both the original dataset and the prediction from the original dataset raster leaves only the walls that have been 'removed'.

Because of the discrepancies in wall placement between the original data and the prediction raster, it was not possible to use raster algebra (original – prediction = result). Instead, it was necessary to create a proximity raster and use this raster to delete any wall within 25 pixels (10 m) distance of the wall. This lead to the final 'found' and 'removed' walls being represented as shorter than they were in reality. However, it gave the clearest and most straightforward to interpret results. Lastly, a field visit was conducted to the study area to validate some of the predictions obtained.

4. Results

The methodology applied produced overall positive results, with the model outputting a general identification of the stone walls in Ærø, as well as new structures classified as potentially new walls. A description of the results is done in a quantitative and qualitative assessment, performed below.

4.1. Quantitative assessment

The loss results per epoch is shown on Figure 9, averaged from after five iterations of the model. The average run had 21 epochs with early stopping (set to run for 50 epochs), which monitored the validation loss (the loss calculated for the validation data) after reaching a maximum of five epochs without improvement on its value.

The evaluation of the test data shows an MSE value of 7,29 (Figure 9), indicating the summed error for all pixels that compose one patch. This was useful to compare between the iterations of the model and during the testing of other additional attempts,



Figure 9 (Left) Median loss and validation loss for all the model's final iterations. (right) Summary statistics of the model training, validation and evaluation, test, data.

alterations, and improvements to the model. Since comparable data or previous research is not available, we cannot compare such evaluation results with reference values. These values were used to measure the results from constant optimisation and improvement of the model. The results from the model training with using only the DTM, or DTM and Height above Terrain (HAT) were slightly similar in comparison to the ones obtained with the three layers, where the evaluation loss and MAE had higher values (Table 1).

With more significance in the evaluation metrics' values, the inclusion of absence data improved the results and significantly lowered MSE and MAE values. The absence data added represented the areas where stone walls are not supposed to be present (mainly representing urban areas). The difference between the pixel values predicted and the true pixels representing a stone wall for the validation area also changed with the different data layers used in the prediction model. The addition of the Sobel filter and the HAT increased the similarity on the number of pixels representing stone walls during the prediction on the test data. The final model adds an average of 0.88 for the pixel-wise metric, corresponding to an 88% overall match between the number of pixels predicted true and the and actual true data. Importantly, numbers should be considered with some reservation since they do not provide much knowledge on the accuracy in the prediction of existing structures or new walls; however, they do indicate the differences in results between model runs. More significatively, the pixel-wise metric calculated for the validation area, from the overall prediction of the municipality of Ærø, was 0.93 for the final model, indicating a high prediction performance.

The output of the predictions on the evaluation data (test dataset) shows a clear picture of the ability of the terrain data to identify stone walls on the landscape. Figure 10 shows some examples of predicted stone walls (first row (a)), where images i and iv display a clear detection of stone walls, as confirmed by the ground truth data (rows b and d) and stone walls dataset (row c). In column iii, an example of an unsure result of a stone wall prediction is displayed, where no stone wall exists in the dataset, and the aerial image shows on what might be a ridge, in need of further validation.

	DTM	DTM & HAT	DTM, Sobel & HAT Without Absence Data	DTM, Sobel & HAT
Loss Function (MSE)	6,329	6,414	11,029	7,294
MAE	0,015	0,014	0,024	0,016
RMSE	0,079	0,080	0,105	0,085
Pixel-wise metric	0,86	0,88	0,83	0,88

Table 1. Evaluation metric for different training datasets.

2200 😉 E. F. L. TROTTER ET AL.



Figure 10 Example of patches of stone wall predictions and respective labels and ground truth. Columns (i), (ii) and (iv) are examples of correctly predicted stone walls, and column (iii), a suspiciously predicted stone wall. Horizontally: (a) prediction patches; (b) DTM in hillshade; (c) stone wall dataset (label); (d) aerial image from kortforsyning.Dk (Spring 2020).

4.2. Qualitative assessment and external validation

Before considering using DL, an attempt at applying a raster-based feature extraction method was made by applying the Sobel filter to be able to identify stone wall structures. The results of such an approach demonstrated the difficulty in identifying stone wall structures, and differentiating them from modern walls structures, or other structures such as ditches. It was assumed protected stone wall structures have more complex structures, and do not have a similar signature, shape or form across the landscape.

The results from applying the DL technique previously described show a clearer identification of stone wall-like structures (Figure 11), and presents less false positives structures identified by the model. Further attempts using off-the-shelf GIS techniques did not bear satisfying results, so that the application of a DL-based approach was considered.



Figure 11 (a) example of an area with Sobel filter, in hillshade for better visualization. (b) results of the predictions by the DL model, represented in color. (In yellow: the original stone walls dataset). on (a) it is possible to visualize wall structures being identified by the feature extraction method, clearer on the bottom-right corner, while on (b), the prediction results from the DL model discard on what it seems to be road-side embankment structure.

The final predictions for the municipality of Ærø showed positive results for both the identification of removed segments and walls, and new stone walls. The visualisation of the final predictions suggests that the model can generally identify stone walls corresponding to the stone wall dataset, even in areas with more dense vegetation. However, for walls located in areas of dense forest or wall structures that have very little salience, some errors of false-negative predictions can occur (Figure 12).

For the models trained only on the DTM, or DTM and HAT layers, it is possible to identify the differences in the predictions. For the first one, a lower ability to identify the stone walls was detected, especially in differentiating amongst different types of edges, while for the latter, the level of noise, composed of low pixel values, was considerable (Figure 14). A considerable difference is also visible for the predictions that included absence data, where much of the noise and scattered pixels were removed (Figure 13).



Figure 12 Examples of model predictions: (a) shows an example of the challenge in predicting on dense forested areas; (b) examples of registered as well as unregistered walls predicted.



Figure 13 Difference between prediction with absence data (a) and without (b).



Figure 14 Difference of predictions between using DTM alone (left), DTM and HAT (center), and the final model (DTM, Sobel filter and HAT) (left). in yellow: stone walls dataset.

In order to test the generalisation of the model, external validation was done by predicting for a new dataset. The municipality of Silkeborg in the region of Jutland has a more diversified landscape and has a larger geographical size.

On a visual inspection, the prediction shows an overall good generalisation of the model, where the stone walls are distinguished (Figure 15). New, well-defined stone walls appear on the prediction, suggesting a good application of the model for detecting new potential stone walls in Denmark.

4.3. Post-processing

An initial post-processing method was done to reduce the noise created by the edges of the patches by applying offsets during the prediction. Such a method allowed for the creation of clearer prediction images and eliminated unrelated values. For the visualisation of the predictions, a visual scale was applied, where first values between 0 and 1



Figure 15 Example of external validation predictions in Silkeborg: (a) shows the prediction of registered walls and a possible unregistered wall; (b) the prediction follows the stone walls registry; (c) possible unregistered stone walls are clearly predicted; (d) shows where probable inexistent registered walls that were not predicted as such.

would be considered (corresponding to the reference values of stone walls), but ending up including only values above 0.2. Most of the values below this threshold were shown to be mostly noise values.

The final step of post-processing analysis, which highlighted the 'removed' and 'found' stone walls from the final prediction, showed that around 391 stone walls (Figure 15) were flagged with either segments or entire removed walls (a total of around 37 km). These were primarily located in agriculture areas, where it is possible to identify whole segments of wall removed on crop fields and segments on the edges of walls for perhaps accessibility purposes.

Many new stone walls were found in the prediction for Ærø, where varying sizes and definitions could be identified. Such structures were then filtered by their pixel value and length, where only structures with values above 0.50 were considered, and smaller structures with less than 10 meters were excluded (Figure 16).

A field visit was conducted to the study site to evaluate the results against ground truth. It was possible to verify some of the predictions, namely the removal of stone walls edges, during the onsite visit. An entire wall that was identified as removed by our predictions was also verified and confirmed. Additionally, verification onsite allowed for



Figure 16 Stone walls that were predicted as removed or damaged after post-processing.



Figure 17 Examples of post-processed found walls in Ærø (in orange; in yellow, the original stone walls dataset).

the confirmation of predictions of unregistered stone walls, confirming in at least some of the cases, the existence of similar structures to protected stone walls, indicating the positive output of the prediction results (Figure 17 and Figure 18).

Additionally, prediction errors were also verified, and the validation onsite allowed us to perceive the nature of the predicted stone walls. This included structures identified as unregistered stone walls, which were in fact embankments or ridges (Figure 19).



Figure 18 Onsite photo 1: hidden wall structure, where it is possible to observe the untouched wall primarily composed of stones and rocks. This structure was predicted as a stone wall, and it is not included in the official registry.



Figure 19 Onsite Photo 2: The model indicates a stone wall running parallel to a sealed section of the road. This transpired to be a small earthen embankment. These sections of predicted wall span the length of the study area and indicate a systemic inaccuracy in the prediction.

The final results of the analysis can be seen here displayed in a prototype of a WebGIS visualisation tool. This tool was created using Leaflet (Agafonkin 2021) for JavaScript (Figure 20).

2206 🛞 E. F. L. TROTTER ET AL.

Stonewall Visualization Tool



Figure 20 Example of a prototype for a WebGIS to visualise the results of the predictions.

5. Discussion

Our results suggest that LiDAR-derived digital elevation data can be used to extract terrain features from digital elevation data, as also found by (Guyot et al. 2021; Maxwell, Pourmohammadi and Poyner 2020)) and (Chase et al. 2012). Our first step sought to validate the stone walls dataset by comparing them with the DTM and analyzing their elevation profile. Interestingly, this method proved successful in identifying protected stone walls or segments of walls that no longer are present in the landscape, but are still registered, which was easily validated by inspecting the aerial images and elevation model to confirm the results. This method was applied as feature engineering for the Deep Learning step in the project, preparing the proper target data to be compatible with the model training (He, Zhao and Chu 2021).

The application of Deep Learning techniques, specifically the use of Convoluted Neural Networks on digital elevation data, presented promising results in identifying specific pixels where stone walls are present. Pixel-wise based analysis of the output predictions in the validation area suggested a high level of regression-based accuracy (0.93), and on the overall area, an average of 0.88. The algorithm discovered areas where the existing dataset needed to be updated due to the removal of stone walls and new stone walls. These discrepancies seemed to be correct with the results from the first step analysis, as well as manual verification on location, where a selected number of new predicted stone walls presented similar characteristics as the ones already registered, however, this remains to be confirmed by an expert in stone walls. Additionally, the post-processing of the predictions identified a total of 391 stone walls removed or having segments that were removed/taken down in Ærø.

The use of multiple data sources showed an improvement over using the DTM alone. The model results improved with the inclusion of additional data layers, with the best model being trained on a combination of the DTM, HAT, and DTM (Sobel) layers. The visual inspection showed a decrease in noisy pixels on detected edges that are not stone walls, especially in urban areas. Future studies should look at the effects of including supplementary data such as aerial imagery, historical maps, or the locations of municipality borders, where the presence of walls is most likely. Given the relation of stone walls with vegetation, where some structures have vegetation and even trees located on top, the use of NDVI imagery could potentially yield interesting results. A future study should relate the presence of stone walls with a biodiversity measurement such as the bioscore (Ejrnæs et al. 2018), relevant not only for structures detection, but also for promoting their value in biodiversity conservation in areas of persistent habitat fragmentation due to agricultural development.

The results of this analysis could be compared to a raster-based analysis where the Sobel filter or comparative edge detection techniques could be used. Such an approach was initially considered but dismissed after the first analysis. The difficulty with this approach is that it is far less discriminating, identifying all of the edges present in the image, making it necessary to differentiate the walls from other edge-type objects. For this specific case, it showed that additional spatial context is needed in order to distinguish amongst the different structures. The findings of this study show a significant improvement on the edge-detection method because the algorithm is able to distinguish a stone wall apart from most edge-like structures. However, given the unique nature and context of our dataset, we are not able to relate our findings with other studies approaching the same problem, nevertheless, our findings do support those of Øivind, Cowley and Waldeland (2019) and Maxwell, Pourmohammadi and Poyner (2020), who also explored the value of CNNs for extracting features from digital terrain data.

The dataset itself diverges from classical deep learning problems in that the training data itself is not completely validated. Given the data science adage 'garbage in garbage out', it is paramount that the training data is correctly labeled in all machine learning tasks. In the case of our study, the original dataset from which the study is based is not actually representative of ground truth. The first step of our analysis sought to curb the effects of this issue, by removing as many walls as possible from the dataset that were either absent, or dubious. Additionally, given that the task was to find new walls within the dataset, it is notable that our aim was not to achieve the lowest possible loss value, as a perfect agreement between the prediction and the test data would suggest no new walls, which we knew not to be the case.

The advantage of the analysis of DEM data is that the algorithm is able to identify patterns that are not visible using aerial or satellite spectral imagery. This useful in archeological applications, for example, where vegetation cover can be an issue (Chase et al. 2012).

There are some limitations associated with our study. Firstly, we are limited by the availability of our terrain data. This data is available for the entirety of Denmark, in connection with new aerial LiDAR missions, renewed on a rolling basis, depending on the region (SDFE. 2020). However, the Danish digital elevation model is only released every 5 years. The data used in this study was from 2014, and therefore is unable to detect and map changes that have occurred in the intervening time. The results can also be highly dependent on the quality of the LiDAR-derived data, where point cloud densities can influence the ability to detect and correctly identify small and narrow objects and structures, as mentioned by Angelidis et al. (2017). In order to apply the same analysis on an updated product from a new LiDAR mission, such differences and consequent biases

need to be considered. Additionally, while terrain data is available for Denmark, it is not necessarily available in other locations nor with the same characteristics. Although the application of this study can be generalised to the context of other countries with similar protected structures, such as Sweden, Norway or Scotland, it is dependent on the level of resolution of their available national LiDAR data, a relevant component, as confirmed by Angelidis et al. (2017).

Due to computational time required, it was not possible to further optimise the CNN method for the task of identifying stone walls. Exhaustively testing the model for the best possible hyperparameters and architecture was not practical given time constraints and computer power and was also outside the scope of this study. In either case: 1) all findings would probably require validation by an expert in the field, 2) small increases in model accuracy were unlikely to result in the discovery of new walls, only that the walls extent would be slightly more accurate. If this method were to be utilised for a similar task, it may be worth experimenting with other model architectures such as Mask *R*-CNN, as has been used by Maxwell, Pourmohammadi and Poyner (2020), or ResNet, by Øivind, Cowley and Waldeland (2019).

Some difficulties were encountered in applying the CNN performing a regression task, given the challenge in evaluating the results. Considering that the target data presented a high imbalance between presence and absence data for each patch extracted, a classification into wall/no-wall proved to be less stable while running the same model with a classified output (where the last layer activation function was switched to Sigmoid, and the loss function to binary cross-entropy). Therefore, deep regression was applied for this specific problem, whereby the output pixel values give a float value representing the probability of presence of a stone wall, rather than a presence or absence classification. Nonetheless, the regression-based method produces results which are harder to interpret and compare. Furthermore, to the best of our knowledge, there are no previous studies on automatically identifying stone walls that would enable a comparison of performance or results, where the only reference source is the stone walls dataset itself.

The application of this study is relevant to the municipalities of Denmark, and to other countries that have an interest in verifying the position of historical stone walls, such as England and Ireland. It can contribute to the automatisation of the identification and update of the stone walls' registry, and in this way fulfil the recommendations outlined by the Ministry of Culture in Denmark (Christensen 2020). The development of such tool can come in a shape of a Decision Support System, where each municipality could visualise and apply analysis on their specific dataset, and in this way, contributing to the update of the national registry. A prototype is in development, and will also require the involvement of experts, and feedback of its usability by municipalities. Additionally, it would also be interesting to consider the benefits that citizen science can offer, whereby citizens could update the stone walls registry by providing information in the field.

6. Conclusion

This study demonstrates the use of a CNN Deep Learning model for extracting features from Digital Elevation Data to map stone walls in a study site in Denmark, and in this way, update their registry. We used publicly available data and concentrated on the Danish municipalities of Ærø and Silkeborg (for external validation).There was an overall

agreement of 93% between ground truth and the prediction of stone walls in the validation area using pixel-wise evaluation. Good results were seen using the DTM alone, however, better results were obtained when adding HAT and an additional DTM layer with a Sobel filter applied. Good generalizability was found when externally validating the model on new data, showing good results for either the existent stone walls and predicting new potential ones. The method performed best in open areas; however, positive results were also seen in forested areas, which suggests that this method could be useful in the identification of features that might be challenging to detect, using remote sensing techniques alone. In order to further improve the identification of stone walls, we suggest that the inclusion of a multi-modal dataset could be beneficial to add additional context and improve differentiation. Further improvements can also be made by exploring the methodology and optimisation of the deep learning CNN model. The next steps for this research is the provision of the model outputs in a web-based spatial decision support system that helps municipalities maintain an up-to-date registry of stone wall structures.

Note

1. The code used in this study is openly available in the associated GitHub repository: https://github.com/AnaCMFernandes/stonewalls.

Acknowledgements

The present manuscript is based on the dissertation entitled "Machine Learning for Automatic Detection of Historic Stone Walls Using LiDAR Data", produced by the authors E. F. L. T. and A. C. M. F. (Fernandes and Trotter2021).

Data availability statement

The code used in this study is openly available in the associated GitHub repository (https://github. com/AnaCMFernandes/stonewalls).

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Ana Cristina Mosebo Fernandes (D) http://orcid.org/0000-0002-2328-0928

References

Abadi, M., P. Barham, J. Chen, Z. Chen, A. Davis, and J. Dean 2016. "Tensorflow: A System for Large-Scale Machine Learning." 12th Symposium on Operating Systems Design and Implementation, 265–283. doi:10.5281/zenodo.4724125.

Agafonkin, V. 2021. Leaflet, https://github.com/Leaflet/Leaflet

2210 👄 E. F. L. TROTTER ET AL.

- Angelidis, I., G. Levin, R. Díaz-Varela, and R. Malinowski. 2017. "Assessment of Changes in Formations of Non-Forest Woody Vegetation in Southern Denmark Based on Airborne LiDar." *Environmental Monitoring and Assessment* 189: 437. doi:10.1007/s10661-017-6119-8.
- Carvalho, M., B. Le Saux, P. Trouvé-Peloux, A. Almansa, and F. Champagnat. 2018. "On Regression Losses for Deep Depth Estimation." Icip 2018, 1–5. Athens: HAL. https://hal.archives-ouvertes.fr/ hal-01925321
- Chase, A. F., D. Z. Chase, C. T. Fisher, S. J. Leisz, and J. F. Weishampel. 2012. "Geospatial Revolution in Mesoamerican Archaeology." *Proceedings of the National Academy of Sciences* 109 (32): 12916–12921. doi:10.1073/pnas.1205198109.
- Chen, Z., B. Gao, and B. Devereux. 2017. "State-Of-The-Art: DTM Generation Using Airborne LIDAR Data." Sensors 17 (1): 150. doi:10.3390/s17010150.
- Chollet, F. others . 2015. Keras. GitHub.
- Christensen, H. A. 2020. Nyt forsøg: Kulturminister vil kortlægge beskyttede sten- og jorddiger, December 1. https://landbrugsavisen.dk/nyt-fors%C3%B8g-kulturminister-vil-kortl%C3%A6ggebeskyttede-sten-og-jorddiger
- Ejrnæs, R., J. Bladt, J. Moeslund, A. Kirstine Brunbjerg, and G. Brian Groom. 2018. "Biodiversitetskortets bioscore." February 1. Accessed May 2021. https://dce.au.dk/fileadmin/ dce.au.dk/Udgivelser/Notater_2018/Biodiversitetskortets_bioscore.pdf
- Fernandes, A. C. M. and E. F. L. Trotter. 2021. "Machine learning for automatic detection of historic stone walls using LiDAR data." Final Dissertation for the Masters In Geoinformatics, Aalborg University.
- Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep Learning*. Cambridge: MIT Press. http://www. deeplearningbook.org
- Guyot, A., M. Lennon, T. Lorho, and L. Hubert-Moy. 2021. "Combined Detection and Segmentation of Archeological Structures from LiDar Data Using a Deep Learning Approach." *Journal of Computer Applications in Archaeology* 4 (1): 1–19. doi:10.5334/jcaa.64.
- He, X., K. Zhao, and X. Chu. 2021. "AutoMI: A Survey of the State-Of-The-Art." *Knowledge-Based Systems* 212: 1–27. doi:10.1016/j.knosys.2020.106622.
- Kingma, D. P. and J. Ba. 2015. "Adam: A Method for Stochastic Optimization." Iclr, 1–13. San Diego, Ithaca, NY. arXiv.org. https://arxiv.org/pdf/1412.6980.pdf
- Kulturministeriet. 2004. Bekendtgørelse om beskyttede sten- og jorddiger og lignende, BEK nr 1029 af 21/10/2004, October 21. https://www.retsinformation.dk/eli/lta/2004/1029
- Kulturministeriet, K. 2009. Vejledning Om Beskyttede Sten- Og Jorddiger. Copenhagen: Kulturarvsstyrelsen. www.kulturarv.dk.
- Kulturstyrelsen, Kultur Ministeriet Slots og. 2020. *Sten- og Jorddiger Konkret udpegning*. July 30. Accessed March 2021. https://slks.dk/omraader/kulturarv/arkaeologi-fortidsminder-og-diger /sten-og-jorddiger/konkret-udpegning/
- Kulturstyrelsen, Slots- og. 2011. Beskyttede Sten- Og Jorddiger. Dataset. Miljøportal. https://area linformation.miljoeportal.dk/html5/index.html?viewer=distribution
- Lathuilière, S., P. Mesejo, X. Alameda-Pineda, and R. Horaud. 2020. "A Comprehensive Analysis of Deep Regression." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42 (9): 2065–2081. doi:10.1109/TPAMI.2019.2910523.
- LeCun, Y., Y. Bengio, and G. Hinton. 2015. "Deep Learning." Nature 521: 436–444. doi:10.1038/ nature14539.
- Maxwell, A. E., P. Pourmohammadi, and J. D. Poyner. 2020. "Mapping the Topographic Features of Mining-Related Valley Fills Using Mask R-CNN Deep Learning and Digital Elevation Data." *Remote* Sens 12 (3): 547. doi:10.3390/rs12030547.
- Maxwell, A. E., M. S. Bester, L. A. Guillen, C. A. Ramezan, D. J. Carpinello, Y. Fan, F. M. Hartley, S. M. Maynard, and J. L. Pyron. 2020. "Semantic Segmentation Deep Learning for Extracting Surface Mine Extents from Historic Topographic Maps." *Remote Sens* 12 (24): 4145. doi:10.3390/ rs12244145.
- Øivind, T. D., D.C. Cowley, and A.U. Waldeland. 2019. "Using Deep Neural Networks on Airborne Laser Scanning data: Results from a Case Study of Semi-automatic Mapping of Archaeological Topography on Arran, Scotland." Archaeological Prospection 26: 165–175. doi:10.1002/arp.1731.

Pedersen, J. 2021. Jord- og stendiger, https://aktiv.dn.dk/sagsarbejde/natur/jord-og-stendiger/

- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, and J. Vanderplas. 2012. "Scikit-Learn: Machine Learning in Python." *Journal of Machine Learning Research* 12: 2825–2830.
- QGIS.org. 2021. QGIS 3.16. Geographic Information System Developers Manual. QGIS Association. https://docs.qgis.org/3.16/en/docs/developers_guide/index.html
- Ramachandran, P., B. Zoph, and Q. V. Le. 2018. "Searching for Activation Functions." Iclr, 1–13. Vancouver, Ithaca, NY: arXiv.org. https://arxiv.org/abs/1710.05941
- Ronneberger, O., P. Fischer, and T. Brox. 2015. "U-Net: Convolutional Networks for Biomedical Image Segmentation." International Conference on Medical Image Computing and Computer-Assisted Intervention, Munich: Springer, Cham. 234–241. doi:10.1007/978-3-319-24574-4_28.
- Sambasivan, N., S. Kapania, H. Highfill, D. Akrong, P. Paritosh, and L. M. Aroyo. 2021. "Everyone Wants to Do the Model Work, Not the Data Work": Data Cascades in High-Stakes AI." Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, Yokohama: Association for Computing Machinery. 1–15. doi:10.1145/3411764.3445518.
- SDFE. 1977-1992. *DTK/4-Cm Kort (1977-1992)*, SDFE, Kortforsyning. https://download.kortforsynin gen.dk/content/dtk4-cm-kort-trykt-1977-1992
- SDFE. 2012. Copernicus Land Corine Land Cover DNK (CLC). *Dataset*. Styrelsen for Dataforsyning og Effektivisering. https://www.geodata-info.dk/srv/eng/catalog.search#/metadata/a05b11b1-50f6-42b3-9796-5d80c2ed2434
- SDFE. 2014. DHM/Overflade. Dataset. SDFE, SDFE Styrelsen for Dataforsyning og Effektivisering. https://download.kortforsyningen.dk/content/dhmoverflade-04-m-grid
- SDFE. 2014. *Dhm/terræn Dataset*, SDFE, Styrelsen for Dataforsyning og Effektivisering. https:// datafordeler.dk/dataoversigt/danmarks-hoejdemodel-dhm/dhm-2015-terraen-brugerdefineretgeotiff/
- SDFE. 2020a. Danmarks Højdemodel, SDFE. Accessed April 2021. https://sdfe.dk/hent-data/dan marks-hoejdemodel
- SDFE. 2020. Flyfotos og laserscanning, https://sdfe.dk/saadan-arbejder-vi-med-data/flyfotos-og-laserscanning
- Shorten, C. and T. M. Khoshgoftaar. 2019. "A Survey on Image Data Augmentation for Deep Learning." *Journal of Big Data* 6 (60): 1–48. doi:10.1186/s40537-019-0197-0.
- Vincent, O. R. and O. Folorunso. 2009. "A Descriptive Algorithm for Sobel Image Edge Detection." Proceedings of Informing Science & IT Education Conference, Macon, GA, USA: Informing Science Institute. 97–106.
- Virtanen, P., R. Gommers, T. E. Oliphant, et al. 2020. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." *Nature Methods* 17:261–272. doi:10.1038/s41592-019-0686-2.