



# Identifying wetland areas in historical maps using deep convolutional neural networks

Niclas Ståhl<sup>a,b,\*</sup>, Lisa Weimann<sup>c</sup>

<sup>a</sup> Jönköping Artificial Intelligence Lab, Jönköping University, Sweden

<sup>b</sup> Skövde Artificial Intelligence Lab, University of Skövde, Kanikegränd 3, 541 34 Skövde, Sweden

<sup>c</sup> County Administrative Board of Jönköping, Sweden

## ARTICLE INFO

### Keywords:

Analysis of historical maps  
Convolutional neural networks  
Wetland management  
Wetland restoration

## ABSTRACT

The local environment and land usages have changed a lot during the past one hundred years. Historical documents and materials are crucial in understanding and following these changes. Historical documents are, therefore, an important piece in the understanding of the impact and consequences of land usage change. This, in turn, is important in the search of restoration projects that can be conducted to turn and reduce harmful and unsustainable effects originating from changes in the land-usage.

This work extracts information on the historical location and geographical distribution of wetlands, from hand-drawn maps. This is achieved by using *deep learning* (DL), and more specifically a *convolutional neural network* (CNN). The CNN model is trained on a manually pre-labelled dataset on historical wetlands in the area of Jönköping county in Sweden. These are all extracted from the historical map called “Generalstabskartan”.

The presented CNN performs well and achieves a  $F_1$ -score of 0.886 when evaluated using a 10-fold cross validation over the data. The trained models are additionally used to generate a GIS layer of the presumable historical geographical distribution of wetlands for the area that is depicted in the southern collection in Generalstabskartan, which covers the southern half of Sweden. This GIS layer is released as an open resource and can be freely used.

To summarise, the presented results show that CNNs can be a useful tool in the extraction and digitalisation of non-textual information in historical documents, such as historical maps. A modern GIS material that can be used to further understand the past land-usage change is produced within this research. Previously, no material of this detail and extent have been available, due to the large effort needed to manually create such. However, with the presented resource better quantifications and estimations of historical wetlands that have been lost can be made.

## 1. Introduction

Historical maps hold crucial information about the landscape of the past, which is an important part in understanding ecological changes over time (Saar et al., 2012). Older historical maps are drawn by hand without any modern systems to aid, making the layout not fully consistent. It is therefore a time consuming challenge to extract desired information from them. The conventional approach to extract such information is through manual annotation, with the help of various GIS-software. This labour intensive approach causes most current studies of historical landscapes and ecologies to be limited in size and only focusing on smaller areas or regions of particular interest, such as the study by Cousins (2009).

In some cases, automatic extraction of certain certain land covers can be done based on the colouring (Herrault et al., 2013). The tool Hist-MapR (Auffret et al., 2017) is an example of software that has been proven useful for such automatic extraction. However, the fading of colour and the yellowing of old paper is a disadvantage when analysing historical documents. There are also historical maps that are drawn, or digitalised, in black and white and hence do not carry any colouring information. The different land cover, therefore, need to be extracted by methods that analyses the different textures in the map, or discover different land cover implicitly, by analysing the surrounding landscape.

In this paper, we show how artificial intelligence (AI) can take advantage of previous manual annotation efforts that have been conducted. More specifically, it is shown how a *convolutional neural network*

\* Corresponding author at: Jönköping Artificial Intelligence Lab, Jönköping University, Sweden.

E-mail address: [niclas.stahl@ju.se](mailto:niclas.stahl@ju.se) (N. Ståhl).

<https://doi.org/10.1016/j.ecolinf.2022.101557>

Received 1 September 2021; Received in revised form 8 January 2022; Accepted 8 January 2022

Available online 12 January 2022

1574-9541/© 2022 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

(CNN) can be trained using annotated data and thereafter be used to automatically detect areas of specific land cover in an old hand-drawn map. As a proof of concept we show how a CNN can be trained to detect wetlands in the around one hundred years old *Generalstabskartan* Swedish Mapping, Cadastral and Land Registration Authority, which depicts the terrain and land-usage in Sweden. The presented method is trained and evaluated on data from one county in Sweden (Jönköping county). Besides this experiment, the model is also applied to historical maps covering the whole southern part of Sweden, with the aim of generating an overview of historical wetland areas. This overview can be used to quantify the historical wetland coverage. Such quantification enables analysis of the substantial loss of wetland area on a large scale. Hence such material is valuable when deciding where and how wetland restoration projects can be conducted. The result of the presented analysis, together with the source code for the method, is therefore released for public use.

CNNs are known to perform well at partitioning images into different segments, based on their content (Minaee et al., 2021). Approaches similar to the one in the presented work have been used to study historical maps. For example, Saeedimoghaddam and Stepinski (2020) use a CNN to detect road intersections in historical maps provided by the United States Geological Survey (USGS). Another study that uses DL to extract information from historical maps is Weinman et al. (2019) who find and transcribe text from historical maps. There are also some studies on more recent historical maps, such as Le Bris et al. (2020) who segments historical orthoimages coupled to digital surface models (DSM) from the 1980s into different land cover.

Besides the analysis of historical maps, there are several applications utilising CNNs to more recent cartographic material. One such application is the detection and segmentation of different land cover from multi-spectral remote sensing images (Huang et al., 2018). Beside detecting and classifying larger areas, CNNs have been used to detect very specific objects in remote sensing images and they can be applied with enough granularity to provide information about different species of trees (Branson et al., 2018). If more granularity than satellite images can provide is needed, drones can be used to photograph an area. The material, which is collected by the drone, can be analysed with the help of CNNs, and important objects can be detected and extracted. Such approaches have been used to analyse and assess the ecological status of areas and also to detect and track different species Gray et al. (2019a, 2019b).

All these cases highlight that CNNs can be useful in ecological applications that require visual analysis. In this work we show that a CNN based method can be applied to historical maps in order to extract information concerning the occurrence of wetlands. Hence, the presented method can be used to minimise the manual labour that is required for analysing such data. Using data from one region in southern Sweden, the County of Jönköping, we show that the method achieves both an high average precision, recall and  $F_1$ -score.

## 2. Materials and methods

### 2.1. Geographical area of analysis

The presented analysis is conducted on the southern part of Sweden. The limitation to just analyse the southern part arises from the earlier choice of drawing the southern part in a different scale than the northern part. Furthermore, the model is trained on only one of the regions that the material covered. The region that is selected is the Jönköping region, where pre-labelled data already existed. In addition, this region covers several different nature types and a significant part of the region has historically been covered by wetlands.

### 2.2. Convolutional neural networks

The idea behind *convolutional neural networks* (CNNs) was first

presented by Fukushima and Miyake (1982) but the big breakthrough came some years later when Krizhevsky et al. (2012) won the Imagenet competition, which is a competition focusing on object recognition in images. CNNs are a special type of artificial neural networks and are inspired by the receptive fields of the human visual system. The core strength of CNNs are their shift invariant property that allow them to detect a specific pattern in grid-like topologies, for example in an image, independently of the position of the said pattern.

A CNN consists of several layers where, in each layer, a linear kernel is applied to a local area of the previous layer and by this produces a set of new latent features, as depicted in Fig. 1. Directly after the kernel a non-linear function is applied, in the presented research this is a leaky Rectified Linear Unit (ReLU) (Xu et al., 2015). The kernel is repeatedly shifted over, and applied to, the previous representation creating a new grid-structure. An example of the whole process, when a CNN model is used to estimate the probability of a pixel in a historical map to be part of a wetland, is shown in Fig. 1.

In the research literature it is common to include pooling layers after some of the activation functions. Such layers aggregate several neighbouring features into a single representation. Hence, condensing the information of a larger area into a single representation. A typical choice of pooling function is the max or the *sum* function.


### 2.3. Approach

As a proof of concept for the usefulness of deep learning, and more specifically CNNs, in the analysis of historical maps, this paper presents a case study where a historical map of one county in Sweden is analysed. A description of the map is presented in Section 2.4. The CNN method presented in this paper is a full convolutional network, having 7 convolutional layers but no pooling layers. The full configuration of the network is presented in Appendix A. The lack of pooling layers let the input signal flow directly from the input to the output, where the value of the pixel is determined. Furthermore, no padding is added to the image, resulting in the loss of pixels close to the border of the image. However, this is handled when the image segments are extracted from the full map, so that the overlap between the extracted images is large enough to make sure that no area of the map is missed out.

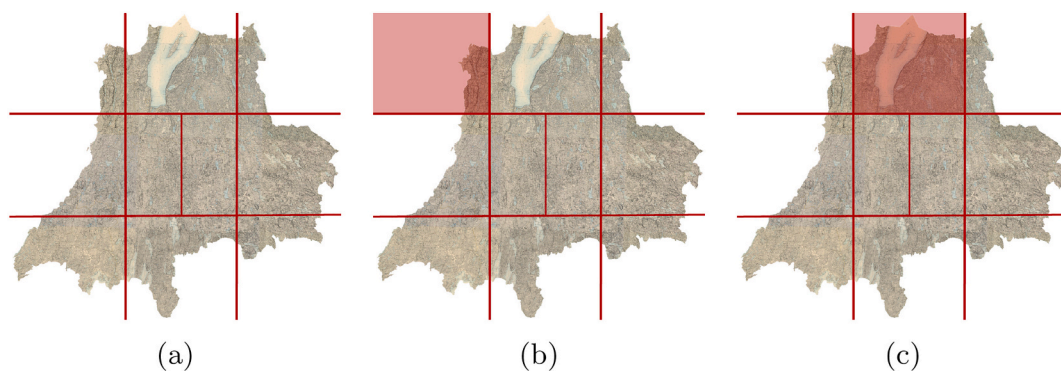
A 10-fold cross validation is performed in order for the result to be generalizable for the remaining maps, to which the CNN is also applied. To create the 10 different sets for the cross validation, we split the map by placing a  $3 \times 3$  grid over the map. The region that is studied is not shaped as a square, and the central cell contains more area than the other 8. This cell is, therefore, split into two cells making it 10 sets in total. The division of the different sets are shown in Fig. 2. During the training of the CNN 9 of these sets are used for the training and the final one is used for evaluation. A challenge to the CNN is that the terrain differs in the different areas, as well as the style of the maps, and thus splitting the dataset in this way would give a good hint on the capability of the CNN to generalise. Among the samples that are used for the training 20% is used as a validation set to prevent the method from overfitting. These samples are selected randomly from all areas that are used in the training set, and are never used to fit the model. The model is trained for 150 epochs and with a batch size of 128. Dropout (Srivastava et al., 2014) with a rate of 0.3 is used during the training to make it more stable. Furthermore, ADAM optimisation (Kingma and Ba, 2014) with a learning rate of 0.0001 is used to find optimal weights in the neural network in order to minimise the cross entropy loss between the network's predictions and the pre-labelled data.

### 2.4. Data pre-processing

As mentioned in the previous sections, the data is first split up into several larger blocks, depending on coordinates, with the purpose to cross validate the model. These blocks are then split into many smaller areas of  $80 \times 80$  pixels, due to limitations in the available amount of



**Fig. 1.** A CNN with four layers. A single pixel is highlighted with a dark blue colour in each of the layers. The values of all feature channels in this pixel are derived from the areas marked in green in the previous layers. In the very first layer the feature channels correspond to the RGB channels of the image, depicting the historical map. The right side of the figure shows the kernel being shifted one step to the right, compared to the left side. This shift results in the prediction of one pixel to the left in the final layer. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



**Fig. 2.** Cross validation of the model. The model is cross validated in such a way that the analysed area is split into 10 sub-areas. These sub-areas varies in terrain type and hence, let us validate the generalisation behaviour of the model to areas with slightly different.

memory. These splits are conducted in an iterative manner so the smaller areas are side by side to each other. In addition, a padding of 27 pixels is added as a frame around the area used as input in order to counter the size reduction that occurs within the CNN. This process creates 41,601 smaller segments, which can be viewed as small images, that are all handled independently by the model.

## 2.5. Data post-processing

Some post-processing is required to transform the result of the CNN into an easy accessible GIS-resource. This is primarily done to produce and refine the material covering southern Sweden, as well as making it easily accessible for further analyses. This process consists of several steps. In the first step, the pixel predictions from the CNN are rounded,

so all predictions with predicted value larger than 0.5 are considered as wetlands and all predictions below are non-wetlands. This creates a raster over the whole map, where each pixel is either deemed to be part of a wetland or not. The next step is to convert this raster representation into a vector representation, to enable further analyses. This conversion is also conducted to minimise storage space and making it easier to distribute.

In the final step, smaller wetlands, which are likely to arise due to noise and oddities in the map, are removed. To achieve this, all wetlands with a total are of less than 1000m<sup>2</sup> are removed.

## 2.6. Software and hardware

All code that are used to produce the results in this paper is released

as open source and is available at Github<sup>1</sup>. The model and all supporting programs are written in Python 3.7 and the model is implemented using PyTorch (Paszke et al., 2019). Furthermore, Rasterio (Gillies et al., 2013) is used for the alignment of the map and the areas that are annotated as wetlands as well as for the rasterisation of the input. Finally, Q-GIS and OpenStreetMap (OpenStreetMap contributors, 2017) are used to visualise the output and generate the output shown in Fig. 5.

### 3. Results

The presented method acquired a  $F_1$ -score, measured over all folds, of 0.886 where the precision of the model is 0.871 and the recall is 0.901. The distribution over each of the different folds, of these three metrics, are shown in Fig. 3. To further dissect how the model functions a smaller excerpt of the map and how the model classifies the different areas are shown in comparison with the annotated areas, in Fig. 4. The total area that was covered by wetlands in the Jönköping region, when the studied maps were drawn, is estimated by the CNN to be  $1.805 \times 10^9 \text{ m}^2$ . This is an overestimation of the wetland area by 0.3% compared to the area that is annotated by humans, which is  $1.8 \times 10^9 \text{ m}^2$ . When the same model is applied to the historical maps that cover the whole southern Sweden, the result of this is shown in Fig. 5, the model estimates the total wetland area in the analysed area to have been  $1.96 \times 10^9 \text{ m}^2$  at the time that the map was drawn. This can be compared to the modern-day wetland coverage in the analysed area of  $7.82 \times 10^9 \text{ m}^2$ , which is provided by the Swedish Mapping, Cadastral and Land Registration Authority. The results shown in Fig. 5 is published in Geodatakatalogen,<sup>2</sup> a portal for sharing data that is hosted by the Swedish County Administrative Boards, and are available at no cost. Furthermore, it can be seen that areas that are known to be densely covered by wetlands today are also the areas which are most densely covered by wetlands in the models predictions, based on the historical map.




Fig. 3. The distribution of precision, recall and  $F_1$ -score over the ten different folds.

<sup>1</sup> <https://github.com/stan-his/GSK-NET>

<sup>2</sup> <https://ext-geodatakatalog.lansstyrelsen.se/GeodataKatalogen/GetMetaDataById?id=4a0e6d1c-abf5-4a30-8819-c4d8e1253bc7&showmetadataview>

Another trend that can be spotted in the detected areas of wetland in the map is that areas that have historically been subject to agriculture contain few wetlands and have already been drained when the historical maps that are analysed were drawn.

### 4. Discussion

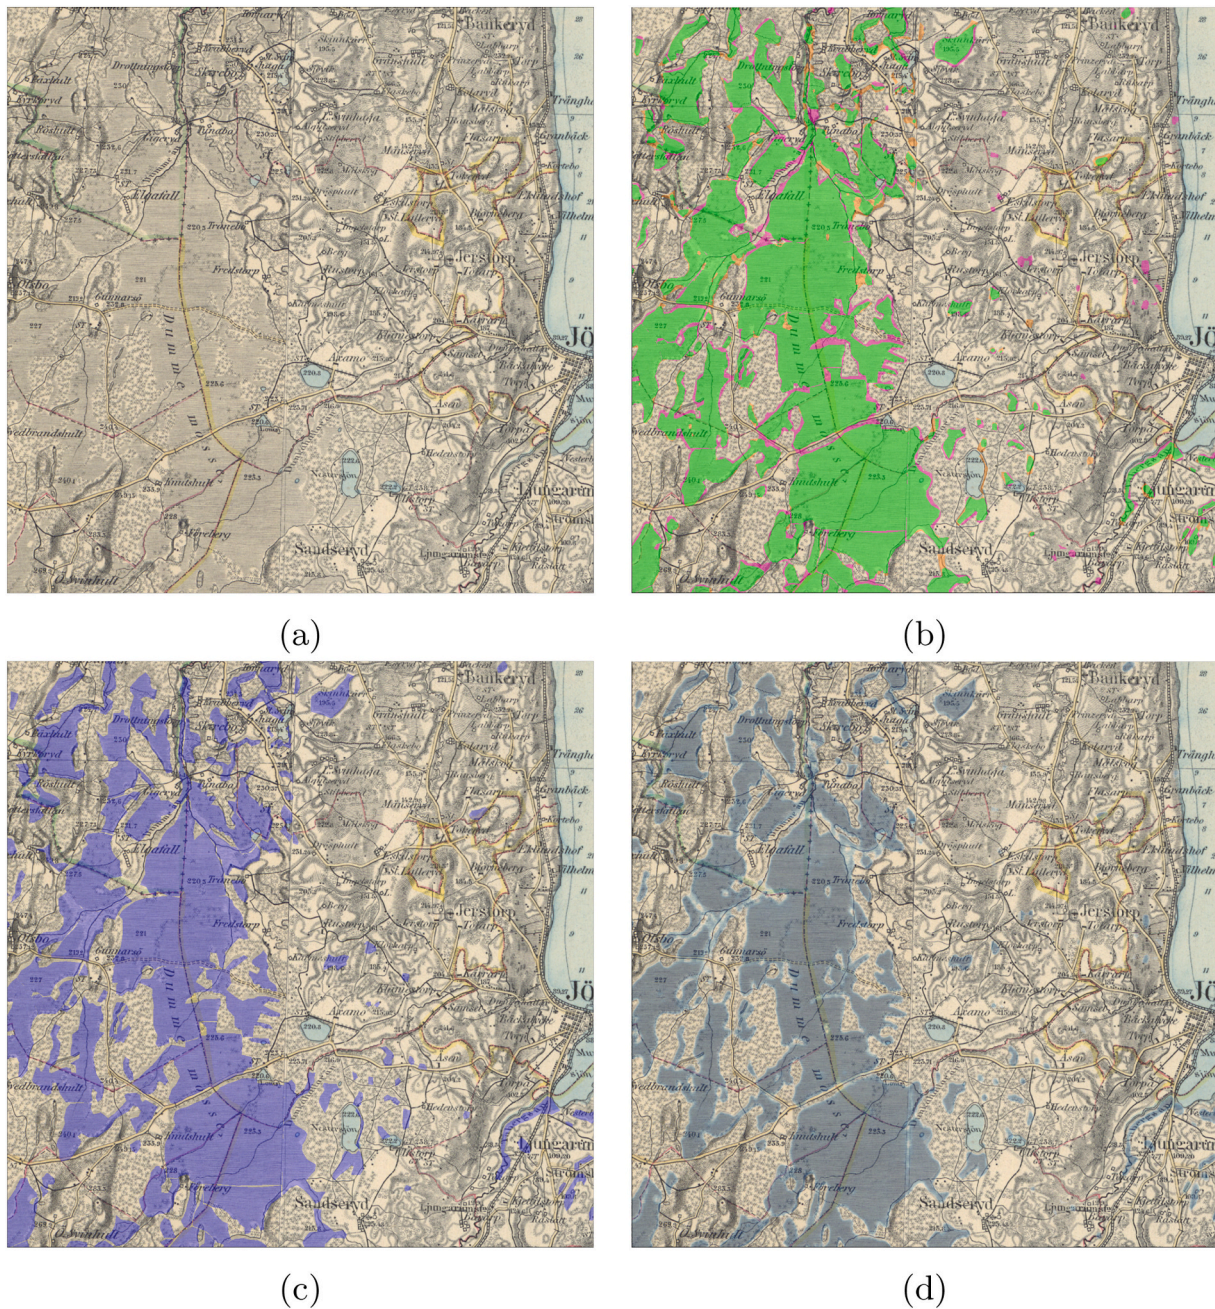
Different AI methods are, more and more, used to reduce the workload that is currently needed to analyse different types material to form a basis for future decisions. This paper explores how a CNNs can be used to detect and segment different land-usages from cartographic material. Several works, which analyse historical cartographic material using CNNs, has been presented earlier with promising results, for example Le Bris et al. (2020) and Huang et al. (2018). However, most of these works are focused on modern multispectral images, most often extracted from remote sensing imagery, which hold less noise and have a coherent standard. It is, however, shown in this paper that this type of approach that can be used, with promising performance, even for historical maps that are full of particularities and only follow a few standards.

The lack of pre-labelled material of high quality, which can be used in the training of supervised models, is a major bottleneck for full scale digitalisation of historical maps. In the presented case, it is shown that data from a single region, covering 173,718 separate wetlands, where the smallest wetland is  $3527 \text{ m}^2$ , is sufficient to get a model with acceptable performance. These wetlands are only selected due to their availability and no formal sensitivity study about the amount of data, or its internal variability, needed to get sufficient performance is conducted within this study. Due to the great heterogeneity of different historical cartographic maps any sensitivity analysis would be dependent on the studied nature type as well as the variability of the representations within the map. Wetlands may be on the easier end of nature types to detect, due to the simple texture, but, there is a risk to confuse the texture with the texture representing hills. Furthermore, the examples that are used, both in the training and the testing, comes from a specific region. Hence, in order to not be biased, a future sensitivity analysis should be done over a multitude of tasks and using samples that are uniformly sampled over the whole area where predictions are expected to be performed.

Since the collection and annotation of data for the training of models, as the one in this paper, is a labour-intensive process, it would be valuable to perform such profound investigations and quantify the amount of data that may be needed. Another way forward, which avoids the labour intense labelling, would be to look for already annotated data that have been used for other purposes and then use that data to build AI models. One technique, which needs to be further investigated, in order to facilitate the collected data better, is to generate additional synthetic samples from the annotated dataset, using generative adversarial models, such as in the presented work by Fang et al. (2019) and Li (2019). In such approaches, several new samples are created in a realistic way, adding more variation and more samples to the original dataset. Hence, a small dataset can be extended with samples that are realistic augmentations of the original dataset (Shorten and Khoshgof-taar, 2019). However, even generative methods requires some annotated data to get started and it is uncertain if artefacts from the generative process will be kept in the generated data and how these artefacts will be expressed.

The presented method only considers the historical maps and no connection to the modern landscape is present. This may cause the generated information concerning the prevalence of wetlands to suffer from reification errors as well as the preservation of artefacts and other flaws from the original map. A future avenue for this research would, therefore, be to couple the generated information together with modern information, of much higher quality, such as soil and elevation data, following a similar approach to Le Bris et al. (2020) where a CNN model with multiple inputs is utilised. Hence, the model would be provided





**Fig. 4.** An excerpt of the map is shown in (a). The areas that are annotated as wetlands, produced by a human, is coloured with blue and shown in (c). The corresponding annotation that is produced by a CNN, which has not seen this part of the map during training, is shown in (d). The similarities and differences between these two annotations are shown in (b). Here the areas for which the human and the CNN annotations are the same are displayed in green. The areas where the CNN annotate the land as wetland but the human did not (false positives) is displayed in pink. Finally, the areas that the humans annotated as wetland but, the CNN did not recognise as such (false negatives) are displayed in orange. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

with a modern GIS-representation of the analysed area and compute an abstract representation of it. This representation would then be merged with an abstract representation of the historical map and the final wetland classification output would, thus, depend on both of these.

## 5. Conclusions

The presented research in this paper shows that it is viable to extract environmental information from historical maps with the help of convolutional neural networks. Our results shows that the performance of the CNN is on par with the human annotations and could minimise the

burden of the manual digitalisation of historical maps. The presented model achieves a  $F_1$  score of 0.87 when a 10 fold cross validation is performed on the data. The disagreement between the CNN and the pre-defined annotation can, furthermore, be explained by small disagreements on how the outline borders of the wetlands should be drawn. This is supported by the fact that the agreement on a macroscopic level, where the agreement between the human annotator and the CNN is almost in unison. In this case, the total area estimated by the CNN differed less than 0.3% compared to the area that was marked by human annotators.





**Fig. 5.** The CNN models estimation of wetlands in the southern part of Sweden, based on the map generalstabskartan. The background map is obtained from OpenStreetMap (OpenStreetMap Contributors, 2017).

### Authors' contribution

The general outline and idea behind this research came from a discussion between the two authors, together with Matti Ermold. LW provided useful insights into wetland preservation work and handled the collection of the required material. NS designed and built the software system and performed the analysis. All authors contributed critically to the drafts and gave a final approval for publication.

### Data availability

The authors do not hold the copyright to the cartographical material that is used to train the presented model. However, this material can be downloaded and used for none-profit purposes, without any cost, from the Swedish Mapping, Cadastral and Land Registration Authority

Swedish Mapping, Cadastral and Land Registration Authority (2021).

Both the wetland data, which is pre-labelled by humans, as well as a GIS-layer of the model result is available through the Geodatakatalogen,<sup>3</sup> a portal for sharing data that is hosted by the Swedish County Administrative Boards, and are available at no cost. The full code to reproduce the performed experiment is available to download from: <https://github.com/stan-his/GSK-NET>.

### Declaration of Competing Interest

The authors have no conflicts of interest to state.

<sup>3</sup> <https://ext-geodatakatalog.lansstyrelsen.se/GeodataKatalogen/GetMetaDataById?id=4a0e6d1c-abf5-4a30-8819-c4d8e1253bc7&showmetadataview>

Acknowledgements

First of all, we would like to thank Matti Ermold at the Swedish Environment Protection Agency who proposed the problem and guided us to relevant material and data and supported us during the project. Secondly, we would like to thank Anne-Catrin Almér and Henrik

Lindblom at the County Administrativ Board of Jönköping for the manual annotation of wetlands which form the data on which the model is trained. Thirdly, We would like to thank Martin Axelsson and Robin Hellgren who explored possible models and worked with the targeted problem and data in their bachelor thesis at the University of Skövde.

Appendix A. Network architecture

Table 1  
Parameters for the CNN model that is used.

Hidden layers	= 6
Neurons	= 128, 64, 64, 32, 32, 32
Kernel sizes	= 9 × 9, 9 × 9, 7 × 7, 7 × 7, 5 × 5, 5 × 5
Dropout rate	= 0.3
Optimiser	= Adam
Learning rate	= 0.0001

References

Auffret, A.G., Kimberley, A., Plue, J., Skånes, H., Jakobsson, S., Waldén, E., Wennbom, M., Wood, H., Bullock, J.M., Cousins, S.A., et al., 2017. Histmapr: rapid digitization of historical land-use maps in R. *Methods Ecol. Evol.* 8 (11), 1453–1457.

Branson, S., Wegner, J.D., Hall, D., Lang, N., Schindler, K., Perona, P., 2018. From google maps to a fine-grained catalog of street trees. *ISPRS J. Photogramm. Remote Sens.* 135, 13–30.

Cousins, S.A., 2009. Landscape history and soil properties affect grassland decline and plant species richness in rural landscapes. *Biol. Conserv.* 142 (11), 2752–2758.

Fang, B., Kou, R., Pan, L., Chen, P., 2019. Category-sensitive domain adaptation for land cover mapping in aerial scenes. *Remote Sens.* 11 (22), 2631.

Fukushima, K., Miyake, S., 1982. Neocognitron: a self-organizing neural network model for a mechanism of visual pattern recognition. In: *Competition and Cooperation in Neural Nets*. Springer, pp. 267–285.

Gillies, S., et al., 2013. Rasterio: Geospatial Raster i/o for Python Programmers.

Gray, P.C., Bierlich, K.C., Mantell, S.A., Friedlaender, A.S., Goldbogen, J.A., Johnston, D.W., 2019a. Drones and convolutional neural networks facilitate automated and accurate cetacean species identification and photogrammetry. *Methods Ecol. Evol.* 10 (9), 1490–1500.

Gray, P.C., Fleishman, A.B., Klein, D.J., McKown, M.W., Bezy, V.S., Lohmann, K.J., Johnston, D.W., 2019b. A convolutional neural network for detecting sea turtles in drone imagery. *Methods Ecol. Evol.* 10 (3), 345–355.

Herrault, P.-A., Sheeren, D., Fauvel, M., Paegelow, M., 2013. Automatic extraction of forests from historical maps based on unsupervised classification in the cielab color space. In: *Geographic Information Science at the Heart of Europe*. Springer, pp. 95–112.

Huang, B., Zhao, B., Song, Y., 2018. Urban land-use mapping using a deep convolutional neural network with high spatial resolution multispectral remote sensing imagery. *Remote Sens. Environ.* 214, 73–86.

Kingma, D.P., Ba, J., 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Proces. Syst.* 25, 1097–1105.

Le Bris, A., Giordano, S., Mallet, C., 2020. Cnn semantic segmentation to retrieve past land cover out of historical orthoimages and dsm: first experiments. In: *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2, p. 1013.

Li, Z., 2019. Generating historical maps from online maps. In: *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pp. 610–611.

Minae, S., Boykov, Y.Y., Porikli, F., Plaza, A.J., Kehtarnavaz, N., Terzopoulos, D., 2021. Image segmentation using deep learning: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.*

OpenStreetMap Contributors, 2017. Planet Dump. retrieved from. <https://planet.osm.org>. <https://www.openstreetmap.org>.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. Pytorch: an imperative style, high-performance deep learning library. In: *Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (Eds.), Advances in Neural Information Processing Systems*, 32. Curran Associates, Inc., pp. 8024–8035.

Saar, L., Takkis, K., Pärtel, M., Helm, A., 2012. Which plant traits predict species loss in calcareous grasslands with extinction debt? *Divers. Distrib.* 18 (8), 808–817.

Saeedimoghaddam, M., Stepinski, T.F., 2020. Automatic extraction of road intersection points from usgs historical map series using deep convolutional neural networks. *Int. J. Geogr. Inf. Sci.* 34 (5), 947–968.

Shorten, C., Khoshgoftaar, T.M., 2019. A survey on image data augmentation for deep learning. *J. Big Data* 6 (1), 1–48.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15 (56), 1929–1958.

Swedish Mapping, Cadastral and Land Registration Authority, 2021. Geographical Survey Office Archive.

Weinman, J., Chen, Z., Gafford, B., Gifford, N., Lamsal, A., Niehus-Staab, L., 2019. Deep neural networks for text detection and recognition in historical maps. In: *2019 International Conference on Document Analysis and Recognition (ICDAR)*. IEEE, pp. 902–909.

Xu, B., Wang, N., Chen, T., Li, M., 2015. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*.