Manaaki Whenua
Landcare Research

Mapping difficult wetlands

Prepared for: Ministry for the Environment

May 2024

Mapping difficult wetlands

Contract Report: LC4468

Brent Martin, Jan Schindler, Andrew McMillan, and Olivia Burge

Manaaki Whenua – Landcare Research

Reviewed by:

Stella Belliss
Researcher Remote Sensing
Manaaki Whenua – Landcare Research

Approved for release by:

Gary Houliston
Portfolio Leader – Plant Biodiversity and Biosecurity
Manaaki Whenua – Landcare Research

Contents

Summary

Project and client

Manaaki Whenua – Landcare Research (MWLR) was contracted by the Ministry for the Environment to conduct research into products and methods for mapping two difficult-to-map wetland types: ephemeral wetlands and wetlands in pasture.

Objectives

- Research and develop methods to assist the national mapping of wetlands, particularly:
  - map ephemeral (mainly kettle-hole) wetlands in the central South Island glacial geomorphology area, within the Canterbury region.
  - research and develop methods that assist the mapping of wetlands in pasture.
- Develop robust methods for national-scale mapping.

Methods

We trialled three methods:

- wetness-based mapping of ephemeral wetlands using indices derived from Sentinel-2 satellite and aerial imagery
- detecting ephemeral kettle-hole wetlands in Sentinel-2 satellite imagery using a deep learning model trained on three mapped areas in Canterbury
- detecting pastoral wetlands in Sentinel-2 satellite imagery using a deep learning model trained on the Whangape catchment in Waikato.

Each method was evaluated by visual inspection against Sentinel-2 RGB images, rural aerial imagery, and maps of known wetlands.

Results

- A process-based water classification model was developed that used 10 m Sentinel-2 imagery to generate the most comprehensive and spatially detailed surface water map to date, and to map a variety of water bodies, including lakes with different colours, braided rivers, ephemeral ponds, and irrigation basins. These layers could be further trained to detect specific types of water bodies, such as small, vegetated ponds. A deep learning model was also trained that could detect very small water bodies from 30 cm aerial imagery.
- A deep learning model trained on mapped areas in the Lake Coleridge and Ashburton Lakes areas successfully detected previously mapped, similar wetlands across the central South Island glacial geomorphology (CSIGG) area, but also detected unmapped wetlands in this area from Sentinel-2 10-band imagery and monthly wetness layers derived from Sentinel-2 imagery. Precision appears to be high in alpine areas, but lower in areas dominated by cultivated land, including forestry.

- A deep learning model trained on the extensively mapped Whangape catchment (using the same process as was used to map ephemeral wetlands in the CSIGG) detected pastoral wetlands in this area with good recall and precision. The model appears to generalise well to similar terrain across the north Waikato and Auckland regions but is less accurate where the terrain or ground cover substantially differs, such as in south Waikato, and on flat, cultivated areas in mid-Waikato.
- These models can be used to generate polygons filtered by wetland probability and size, which could be triaged by mappers and then used to indicate the location and approximate size and shape of potential wetlands for manual mapping.

Conclusions

- We generated a prototype layer councils can use to create candidate wetland polygons, which will help with manual inspection and delineation tasks.
- The deep learning-based water detection mask using high-resolution aerial photography adds fine-scale water features to the data set, and can be used as a mask to constrain the frequency layer and reduce false predictions over forests and other land-cover types.
- Both layers are useful beyond wetland applications and proved stable for many different water body types, including braided rivers, streams, ponds, kettle holes, lakes, and coasts.
- The deep learning model trained on mapped wetlands in the Lake Coleridge and Ashburton Lakes areas generated an ephemeral wetland layer for the CSIGG that may be useful for mapping ephemeral wetlands in this area.
- The deep learning model trained on the Whangape catchment appears to be reasonably predictive of potential pastoral wetlands and might be a useful additional tool for finding this wetland type. The model also appeared to transfer quite well to similar terrain in the Auckland region, demonstrating that it might be useful for other parts of the country with similar land cover and terrain.
- These products do not automate the mapping of wetlands, but they can be used by mappers to quickly identify potential wetlands down to sizes of less than 1 ha, which are then manually mapped. We hope that by combining the products with existing imagery and mapping techniques, it may become more feasible and cost-effective to extensively map wetlands.
- Although this research focused on two case studies – ephemeral wetlands in the CSIGG and pastoral wetlands in the Waikato region – the same approaches are potentially applicable to other regions and wetland types.

Recommendations

- Assess the accuracy and usefulness of the products developed in collaboration with council staff, including an assessment of which (if any) apparent false positives (i.e. detected areas that are not kettle hole ephemeral wetlands (in the CSIGG) or pastoral wetlands (in the Waikato region) are useful for detecting and mapping other types of wetlands.

- Extend the water frequency layer to all Sentinel-2 satellite imagery for the period 2016–2024.
- Experiment with turning the water frequency layer into a classification map (permanent water, seasonal wetness, one-off/flooding events).
- Extend the deep learning training areas to include a greater range of land cover and terrain.
- Add further inputs to the deep learning models, including aerial imagery and synthetic-aperture radar (SAR).
- Increase the deep learning training data time series by adding more seasonal images within the year studied, and potentially multi-year imagery, where possible.
- Carry out further tuning of the model training processes, including trialling alternative methods of modelling temporal change.

1 Introduction

Wetlands are important, rare ecosystems that provide a plethora of ecosystem services, including habitat provision for native flora and fauna, nutrient and water recycling, and cultural and socio-economic benefits (Johnson & Gerbeaux 2004). They are mainly threatened by human activities disturbing water levels and damaging vegetation, with over 90% of wetlands in New Zealand lost since European settlement (Ausseil et al. 2011). To protect these critical ecosystems, we need basic information on wetland locations, and accurate mapping of their extents and types.

Regional councils are required by the National Policy Statement for Freshwater Management 2020 (NPS-FM) to map wetlands and assign them a wetland type. However, mapping small and seasonally changing types, such as ephemeral wetlands, is a technological and pragmatic challenge at regional to national scales in New Zealand.

Previously, technical guides were developed to delimit wetland extent, both manually via desktop methods (MfE 2024) and in the field (Clarkson 2014; Fraser et al. 2018; MfE 2021, 2022). Experts with knowledge in ecology, geomorphology, and botany are needed to accurately delineate wetlands and assign their type. This manual process is expensive and time-consuming, so any assistance with the manual mapping process would reduce the resources needed to achieve the mapping required by the NPS-FM.

This research investigates using remote-sensing methods to generate GIS layers that may assist with mapping difficult wetlands. We focus on two types of wetlands: ephemeral wetlands (particularly kettle-hole wetlands; see Johnson 2003), and pasture-dominated wetlands, both of which were identified by stakeholders as problematic to map. A third difficult wetland type – forested wetlands – was excluded owing to the lack of visibility of the wetlands in readily available remote-sensing data, such as satellite and aerial imagery.

2 Background

The primary focus of this research is ephemeral wetlands. These have seasonally changing water levels and may only be partially inundated for a short period of the year, or every few years, making them difficult to observe in imagery.

Similarly, pastoral wetlands are often visually indistinguishable from other vegetation types most of the time, especially in low-resolution imagery such as 10 m Sentinel-2 satellite data. They are, however, often distinguishable by their seasonal variation.

This research focuses on two complementary approaches to detecting wetlands:

- **wetness-based mapping**: temporal changes in wetness are derived from imagery and used by the mapper to indicate potential wetlands
- **image-based mapping**: potential wetlands are detected from imagery, including the wetness layers as additional image bands, and the resulting detections are used to guide the mapping process.

These two approaches are now described.

2.1 Wetness-based mapping

This part of the research focuses on ephemeral wetlands. The main idea is to compress the Sentinel-2 satellite time series from a whole year to just one layer that carries information on how often certain areas were inundated. We apply state-of-the-art techniques in remote sensing and machine learning (including deep learning), analysing both temporal, medium-resolution satellite data and high-resolution aerial photography to generate candidate areas of potential locations and extents of ephemeral wetlands.

We aim to stay as close as possible to empirical observational data and a physical understanding of temporal patterns in ephemeral wetlands, thereby supporting manual inspection and interpretation. The overall workflows are designed to be transparent and reproducible, with opportunities for improvement at later stages, where appropriate.

2.2 Image-based mapping with deep learning

While both pastoral and ephemeral wetlands are often visually indistinguishable from other vegetation types, especially in low-resolution imagery such as 10 m Sentinel-2 satellite data, they are often detectable from their seasonal variation in appearance. We used deep learning to exploit this fact by training a deep learning encoder–decoder model to process a temporal stack of images to detect potential wetlands.

Deep learning encoder–decoder networks translate one image to another, where the second image is generally less complex than the first. The network is trained by repeatedly applying it to pairs of input and output images, and incrementally adjusting the model's parameters to gradually improve overall performance on the task. This approach can be readily applied to wetland mapping by training the encoder–decoder network on pairs of wetland polygons (converted into an image) and remote-sensing imagery such that the model learns to transform the imagery into a wetland map. The model is trained on a subset of the area of interest where wetlands of the target type have been extensively mapped. The model is then applied to the rest of the area.

Deep learning networks typically produce a *semantic* raster, where each pixel represents the class of the corresponding image pixel (e.g. wetland/not wetland). However, the model can also generate a *probability* raster, where each pixel's intensity represents the model's confidence that it is of the target class (e.g. ephemeral wetland). This provides a richer evaluation of the likelihood that an area is a wetland and can be more sensitive to areas where the visual evidence is weak.

As well as learning to perform the mapping task, encoder–decoder networks can potentially improve on the original map if there are errors; this is because these errors contradict the general patterns learned, and so are ignored during training but corrected when the model is run over the same imagery. This approach has the potential to discover additional wetlands within the same area as that used for training, allowing areas that

have been incompletely mapped to be used for training provided the level of error/omission is not too severe.

We explored using deep learning encoder–decoder networks to transform temporal image stacks into probability rasters that suggest candidate ephemeral wetlands. We then extended the approach to include mapping pastoral wetlands, on the assumption that temporal changes occurring in these wetlands and the surrounding pasture could be used to detect them.

3 Objectives

- Research and develop methods to assist the national mapping of wetlands, particularly:
  - map ephemeral (mainly kettle-hole) wetlands in the central South Island glacial geomorphology area, within the Canterbury region
  - research and develop methods that assist the mapping of wetlands in pasture.
- Develop robust methods for national-scale mapping.

4 Methods

4.1 Data preparation

The experiments described used imagery obtained from processing Sentinel-2 satellite data as the main input. In addition, some of the experiments included rural aerial imagery and a digital elevation model (DEM). A common source of ground truth for ephemeral wetlands was applied across all applicable experiments. Each of these data sources is now described.

4.1.1 Wetland polygons

Ephemeral wetland polygons

We used polygons of ephemeral wetlands from in-field observations that had been manually digitised for a previous study on ephemeral wetland mapping using multi-temporal satellite imagery (McMillan & Wiser 2019). These data were collected by Manaaki Whenua – Landcare Research (MWLR) and the Department of Conservation (DOC) in 2015 in Canterbury, New Zealand (Figure 1).

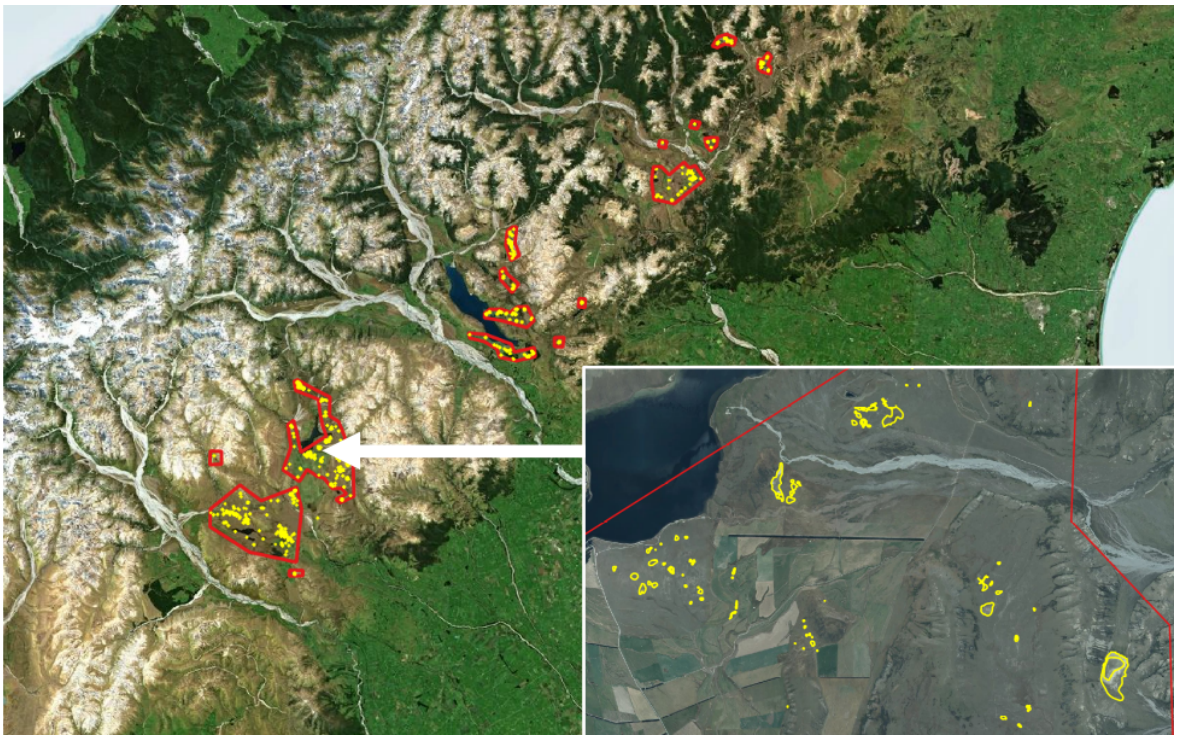

Figure 1. Ground-truth ephemeral wetland polygons (in yellow). The larger red polygons indicate the boundaries of the areas for which ground truthing was available for this project.

Pastoral wetland polygons

Wetlands in the Whangape catchment were mapped in 2021 by MWLR in a desktop exercise, using Waikato rural aerial photos and oblique aerial imagery. These were used as training data. Additional wetland polygons were obtained from Waikato Regional Council's consolidated wetland layer and Auckland Council's ecosystem layer, and these were used for testing purposes. Figure 2 shows the extent of the mapping.

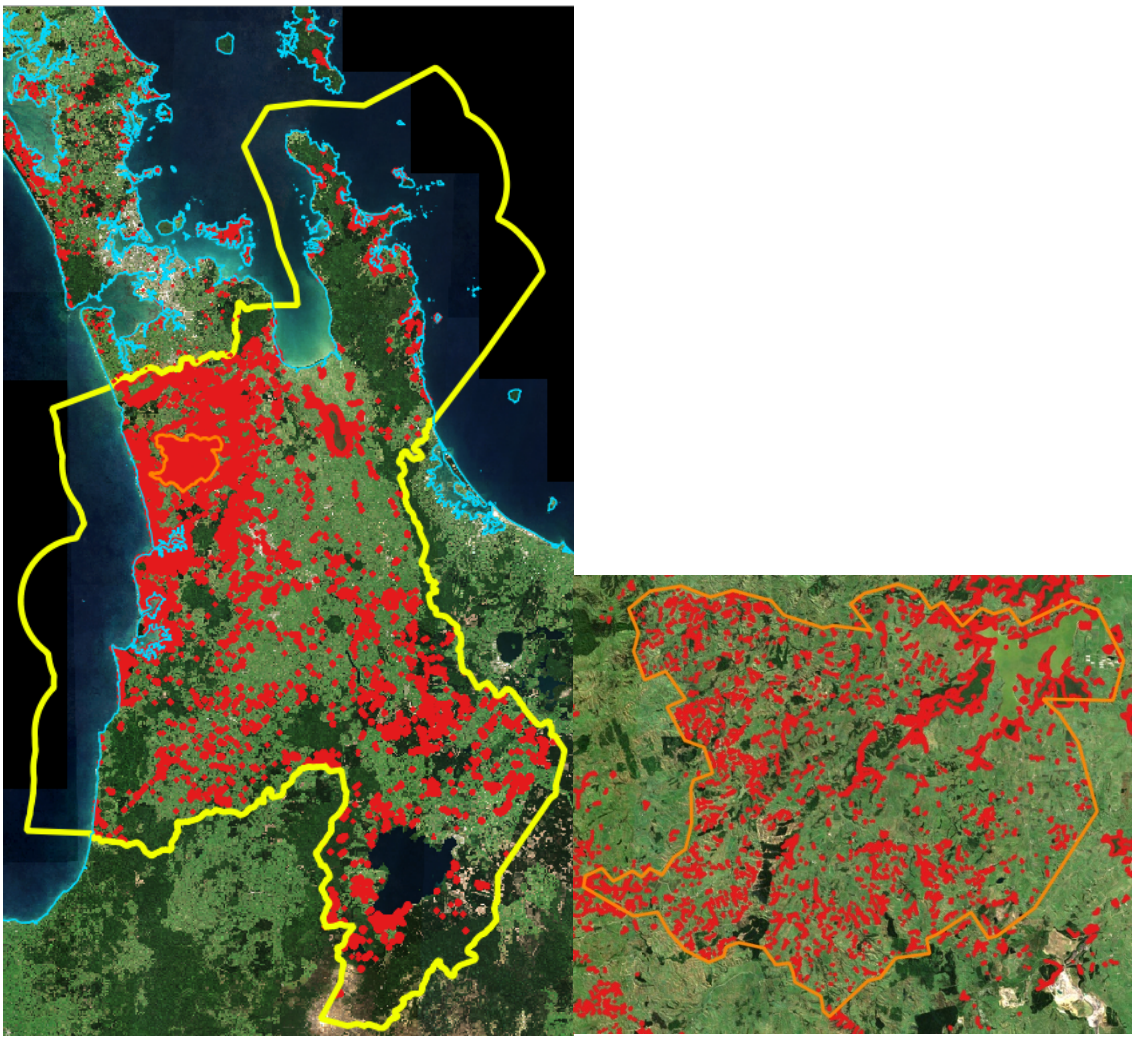

Figure 2. Mapped wetlands in the Waikato (yellow boundary) and Auckland regions (left) and the Whangape catchment (orange boundary, right) over a Sentinel-2 RGB mosaic.

4.1.2 Satellite imagery

Daily satellite imagery for wetness-based detection

We collected all Sentinel-2 (A & B) images (i.e. orbit passes) available for New Zealand from 2020 and 2021 and created reconstituted daily images of complete individual orbits. The ground pixel size of the Sentinel-2 bands in the visible spectral range (B2/B3/B4) and near-infrared range (B8) is all 10 × 10 m (Drusch et al. 2012). The red edge and short-wave infrared bands, which both have 20 m pixels, were sharpened by oversampling to 10 m pixels using local correlation against the 10 m resolution bands (Dymond & Shepherd 2004).

Spectral band data (Table 1) were then processed to standardised reflectance; i.e. the spectral reflectance the surface would have on flat terrain with the sun and satellite in standard positions (nadir view for the satellite and a solar elevation of 50^0), using the algorithms described in Dymond & Shepherd 2004).

Table 1. Sentinel-2 bands used in this study: indicative central wavelengths and bandwidths for Sentinel-2 A & B bands

Band	Short name	Wavelength	Bandwidth	Spatial resolution
2	Blue	492 nm	66 nm	10 m
3	Green	559 nm	36 nm	10 m
4	Red	664 nm	31 nm	10 m
5	Red-edge 1	704 nm	15 nm	20 m
6	Red-edge 2	740 nm	15 nm	20 m
7	Red-edge 3	782 nm	15 nm	20 m
8	NIR wide	832 nm	106 nm	10 m
8A	NIR narrow	864 nm	21 nm	20 m
11	SWIR 1	1,612 nm	92 nm	20 m
12	SWIR 2	2,202 nm	180 nm	20 m

Source: Drusch et al. 2012.

Cloud masks were generated by running a single-image cloud detection method called Fmask (Zhu & Woodcock 2012; Zhu et al. 2015; Qiu et al. 2019), which feeds into a multi-temporal cloud detection algorithm called Tmask (Zhu & Woodcock 2014). Because the resulting pixel-wise cloud mask layer still contained artefacts of false positive cloud and shadow pixels, we then ran an algorithm that also looks at the approximate likely cloud–shadow distances (based on sun azimuths) to remove segments without a corresponding cloud/cloud-shadow partner. In addition, the Sentinel-2 cirrus band was used to mask out very thin, faint, cirrus clouds, which may have an impact on the calculation of spectral vegetation indices.

Finally, the cloud and cloud-shadow segments were smoothed and filled with a generalisation algorithm using multiple morphological operations. (Our Sentinel-2 pre-processing and cloud detection workflow is described in Belliss et al. 2019 and Shepherd et al. 2020.) The number of valid observations per ground pixel varies greatly across the country, mainly due to satellite orbit patterns and overlap regions (more observations), and frequent cloud cover in mountainous areas (fewer observations).

Monthly satellite imagery for image-based detection

For the image-based modelling, images were generated for specific months of the year using a similar process. The mosaics were further processed by replacing all cloud-affected areas with the best possible imagery from alternative tracks from the same month where the affected area was the clearest. First, the cloud (and cloud shadow) affected areas were identified. The cloud-affected area was then cut out of the mosaic and replaced with the same portion of the alternative image. Finally, the mosaic was 'flattened' by minimising the shading effects arising from topology, giving an image that attempts to best represent purely land-cover texture.

4.1.3 Aerial photography

High-resolution aerial imagery at 30 cm pixel resolution with red-green-blue-infrared (RGBi) channels was collected for Greater Wellington[1] and Hawke's Bay.[2] Individual image tiles were merged to form a large mosaic for each region and stored as an 8-bit unsigned Byte-type (0–255 pixel values), four-band raster image in the efficient KEA-format (Bunting & Gillingham 2013).

4.1.4 Slope information

We used elevation information from the New Zealand 8 m digital elevation model (DEM) published by Land Information New Zealand (LINZ) and available online.[3] We re-gridded the DEM to the Sentinel-2 10 × 10 m pixel resolution and calculated a slope map for the country using the GDAL gdaldem tool.[4]

4.2 Wetness-based detection of potential ephemeral wetlands

4.2.1 Time series behaviour of ephemeral wetlands

The main assumption of this research is that periodic changes in ephemeral wetlands are observable in the time series of spectral information from the Sentinel-2 satellite sensor. We investigated the temporal patterns of several ground truth polygons and studied spectral indices that are sensitive to changes in vegetation information and moisture content. One popular index is the Normalized Difference Vegetation Index (NDVI) (Rouse et al. 1974 and Tucker 1979). Its values typically range from -1 to +1 with very high values correlating with dense, healthy vegetation and very low values relating to water or land absent of vegetation. We generated time series plots of NDVI based on the daily images available for the period of 2016-2021 for samples of wetland and non-wetland areas, and visually inspected them for patterns that could be used to inform classification decisions.

4.2.2 National water frequency mapping

The overall workflow process to generate temporal water frequency maps includes five stages: data preparation → training data → classification model training → prediction → summary statistics, which are described in the following sections.

Satellite data preparation

We used our internally established Sentinel-2 archive of daily analysis-ready satellite images for the year 2021. The individual daily images were atmospherically corrected,

[1] https://data.linz.govt.nz/layer/105727-wellington-03m-rural-aerial-photos-2021/

[2] https://data.linz.govt.nz/layer/110466-hawkes-bay-03m-rural-aerial-photos-2021-2022/

[3] https://data.linz.govt.nz/layer/51768-nz-8m-digital-elevation-model-2012/

[4] https://gdal.org/programs/gdaldem.html

cloud- and shadow-masked, snow/ice reduced, and screened for invalid data. We reduced the number of spectral bands to the ones that are sensitive to water absorption (i.e. Sentinel-2 band 2 in the blue spectral region, centred at around 492 nm; band 3 in the green spectral region, centred at around 559 nm; and band 8 in the near-infrared spectral region, centred at around 832 nm). Three-band composite images in this band combination were created for further analysis.

Training data

We manually created two point data sets in a Geographical Information System (GIS) comprising 1,161 points over different water bodies visible in Sentinel-2 imagery and 1,445 points over land and forested areas (Figure 3). We carefully selected Sentinel-2 summer scenes between 2016 and 2023, making sure the water and land point locations were clearly visible. We then extracted spectral signatures of the Sentinel-2 pixel closest to each point location. We decided against averaging spectral values within a buffer region around each location because we plan to apply a pixel-wise classifier at the original 10 × 10 m Sentinel-2 resolution.

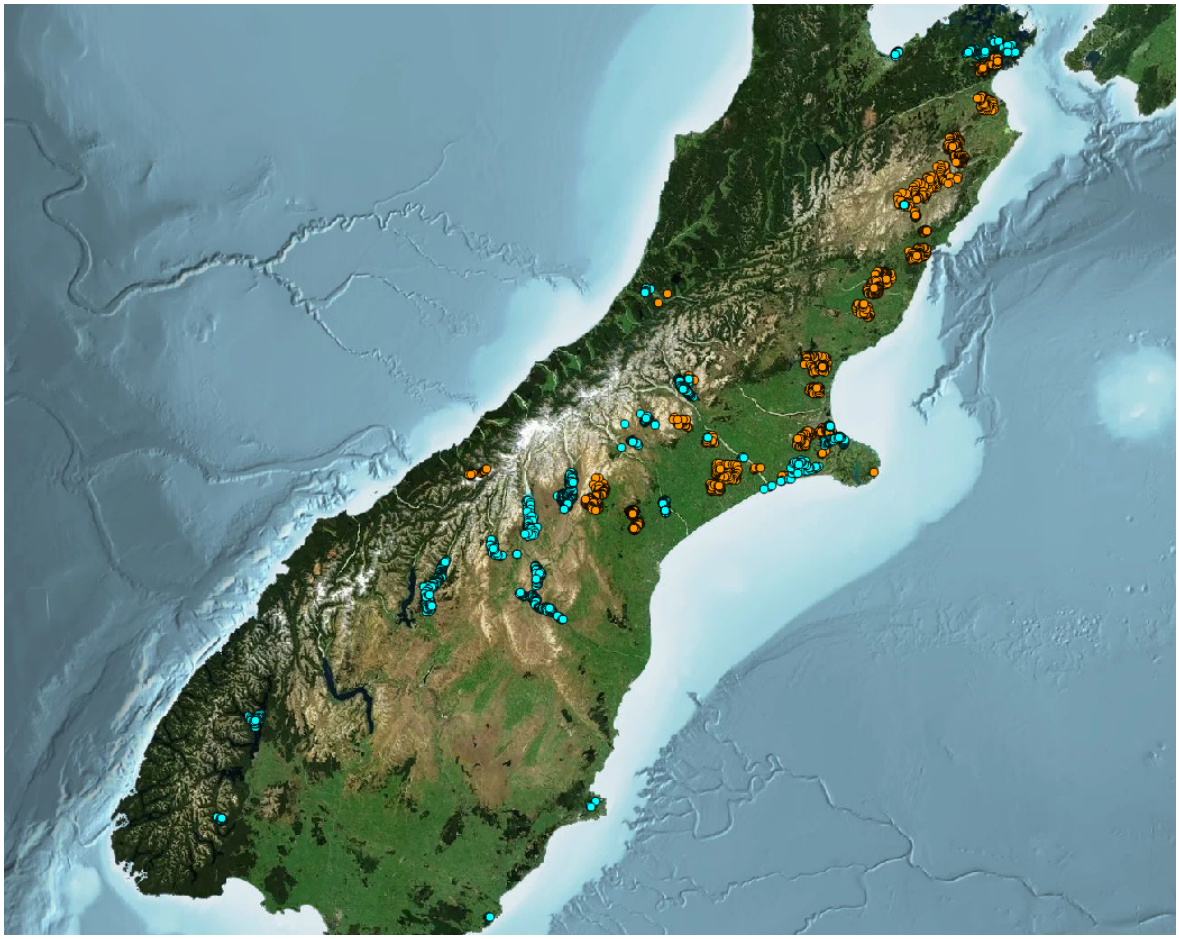

Figure 3. Pseudo ground-truth water (blue) and land (orange) point data manually created in a GIS software.

Water detection model training

We trained a logistic regression classifier model (Cox 1958; Hosmer et al. 2013) using the Python-based Scikit-learn toolkit (Pedregosa et al. 2011). Logistic regression is a simple, fast, supervised machine-learning algorithm ideal for binary classification tasks of large data sets. We trained a water/land model using standard settings:

- penalty: L2
- dual: False
- C: 1
- class_weight: none (all classes are have weight one)
- fit_intercept: true
- intercept_scaling: 1
- max_iter: 1,000
- multi_class: auto
- solver: lbfgs tol: 0.0001.

To speed up computation we used the Intel(R) extension for scikit-learn.[5] The model accuracy of 0.943 (±0.032) was estimated using repeated five-fold cross-validation with stratified random sampling.

Prediction

The trained classification model was then applied to each Sentinel-2 orbit pass (c. 360 per year). Overall, water is detected well, but many false predictions are made in steep terrain, as this is often affected by shadow or snow and can look spectrally similar to some water surfaces. To combat this issue, we used a slope map derived from the LINZ 8 m DEM as mask, with a slope exclusion threshold of ≥10°. Thus, the water body detection is only applied to pixels in relatively flat terrain. Some false positives remained, which can be screened out during the (manual) wetland mapping process.

Water frequency summary statistics

The water classification was applied to each daily image in 2021. We stacked all water masks for that year on top of each other and summed the number of water observations at the pixel level for each month, and then summed the monthly water pixel counts over the year. All computations were performed separately for the South Island and North Island.

4.2.3 Water body detection in high-resolution aerial imagery

The medium resolution of the Sentinel-2 optical sensor at 10 × 10 m ground pixel size restricts the minimum size of water bodies that can be detected. It is therefore expected

[5] https://github.com/intel/scikit-learn-intelex

that small water surface features such as streams or ponds may not be detected using Sentinel-2. To overcome this issue, we ran a parallel experiment to test whether water can be detected directly from high-resolution 30 cm aerial photography using a deep learning encoder–decoder network.

Training data

We hand-delineated 1,870 training polygons covering different types of water bodies, including streams, rivers, wetlands, ponds, lakes, kettle holes, and irrigation systems (Figure 4). We paid close attention to including different water body colours, sizes, and shapes. The training data were collected in the Greater Wellington region (30 training areas with 949 training polygons) and Hawke's Bay (24 training areas with 921 polygons). We also included true negative areas with no water coverage. The aerial RGBi imagery was then split into 512 x 512 pixel square tiles within each training area. A random sampling approach was conducted, splitting the image and labels into 80% training and 20% validation sets.

Figure 4. Manual labels for the deep learning water detection model. Training areas (red boxes) and water body polygons (blue-filled polygons) were created in Greater Wellington and Hawke's Bay.

High-resolution water detection model

We trained a U-Net-type CNN model (Ronneberger 2015) with a Resnet101 encode backbone pretrained on the ImageNet data set, using the PyTorch framework with focal loss and gamma = 2, a batch size of 32, a learning rate of 0.0001, and the following augmentations: horizontal and vertical flip, transpose, rotate90, shift-scale-rotate, and brightness+contrast stretch. We trained the model until the validation loss and accuracy reached their maximum and stabilised.

4.3 Image-based detection of potential wetlands

4.3.1 Detecting kettle-hole ephemeral wetlands using deep learning

We conducted some initial experiments to see whether ephemeral wetlands could be detected at all using a model trained on the limited number of kettle-hole wetlands mapped in Canterbury by DOC and MWLR (McMillan & Wiser 2019) and tested on the same data set.

We trained a U-Net model using the kettle-hole wetland polygons and imagery from the Sentinel-2 10-band, 10 m mosaic for September 2020. The images were tiled into 384 256 × 256 pixel images, of which 50% were randomly selected for training, 10% were used as a validation set to estimate the model's performance during training, and the remaining 40% were withheld for testing.

There are many 'hyperparameters' that can be tuned to maximise the performance of the U-Net model. We conducted a limited grid search over the main parameters and settled on the following settings:

- input image size: 256 × 256 pixels
- number of images in each training batch: 2
- learning rate (determines the size of each update to the model's parameters): 0.0001
- optimiser: ADAM[6]
- length of training runs: up to 100 epochs of 300 batches (each of two images) per epoch; i.e. a total of 100 × 300 × 2 = 60,000 image presentations
- model selection: the model with the lowest loss when tested on the validation set
- augmentation: the following spatial transformations were applied to the training images to reduce the tendency to overfit:
  - randomly flip images horizontally

[6] This optimiser automatically adjusts the size of updates to the model's parameters, based on an approximation of the L1 (mean) and L2 (variance) moments of the error. This strategy makes large changes to the model early in training, when the mean error is large with respect to the variance, then decreases the change size as the error reduces and variance begins to dominate. This automatically throttles back the updates to reduce the likelihood of overfitting.

- randomly flip images vertically
- randomly rotate images by 90˚.

The trained model was used to generate a probability raster representing the likelihood that an area is a kettle-hole wetland. We then assessed the model's success by visually comparing the generated raster image to the mapped wetlands for areas excluded from the model's training set.

Detecting kettle-hole ephemeral wetlands from pairs of Sentinel-2 images

The appearance of the ephemeral wetlands changes considerably depending on the time of year. In McMillan & Wiser 2019, temporal variation in wetness was used to detect ephemeral wetlands by calculating wetness variation indices from 5 years of data. For this experiment we tested whether the variation in wetness between just two images from the wettest and driest months is sufficient to detect kettle-hole wetlands. Of course, the wettest and driest month is variable between years.

We selected the months of March and September after visually inspecting all monthly mosaics for 2020 and 2021 because they consistently showed strong seasonal variation and were mostly free of clouds. Figure 5 shows the seasonal variation in appearance of kettle-hole wetlands in an area south of Lake Heron in Canterbury. In both years the appearance of the wetlands changes significantly between March and September, but there are also striking differences between the 2 years, with the wetlands in 2020 being more strongly vegetated in summer and less flooded in winter compared with the 2021 imagery.

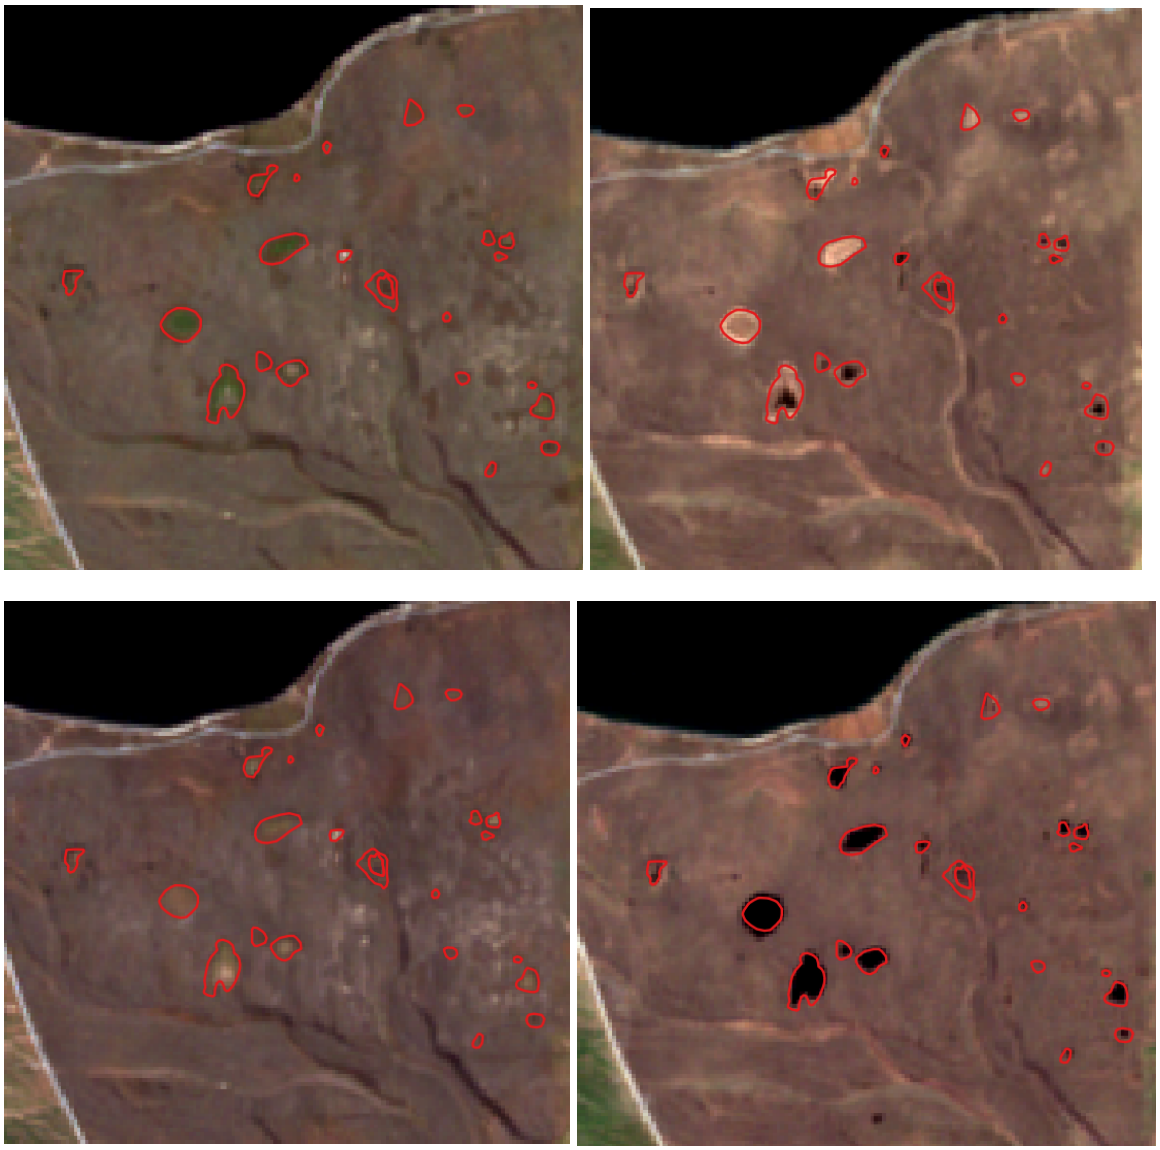

Figure 5. March (left) and September (right) Sentinel-2 images for south of Lake Heron for 2020 (top) and 2021 (bottom). Red polygons are mapped kettle-hole wetlands. Black areas are open water.

We tested this approach by training a U-Net model on stacked pairs of March and September images and then comparing the results to those generated by the previous model which was trained on a single summer image. Figure 6 shows an example of the detection differences in these two: the model that uses only one image (centre) performs poorly, with relatively low detection and a high false positive rate, while the two-image model (right) has a significantly higher detection rate with fewer false positives.

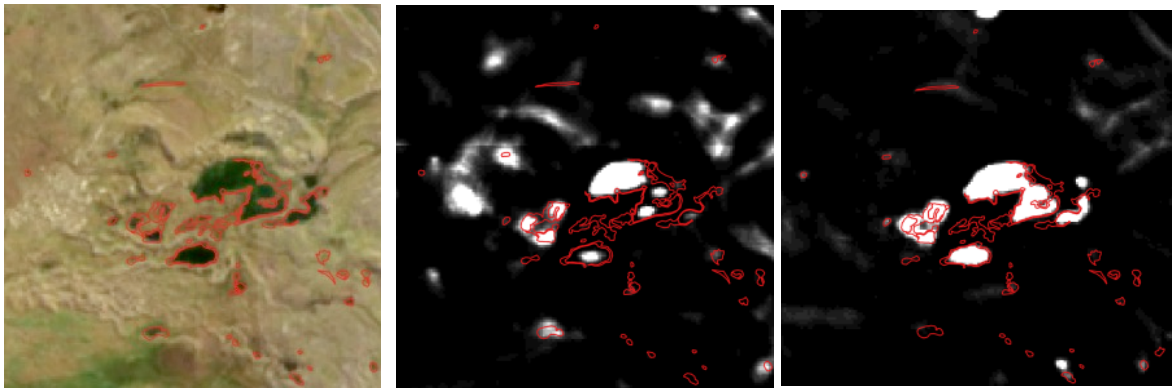

Figure 6. Comparison of single-image and two-image predictions. Bright areas indicate wetland detection; the brighter the area, the higher the probability that it is an ephemeral wetland. The white areas outside the red-outlined polygons are false positives (areas not mapped as ephemeral wetlands).

The two-image model was then refined and used to generate predictions over larger areas to gauge how well the approach performs. In the original McMillan and Wiser study, an area 6 km × 9 km was used to test the models. Figure 7 shows the Sentinel-2 image from March 2021 for this area, and the model's output.

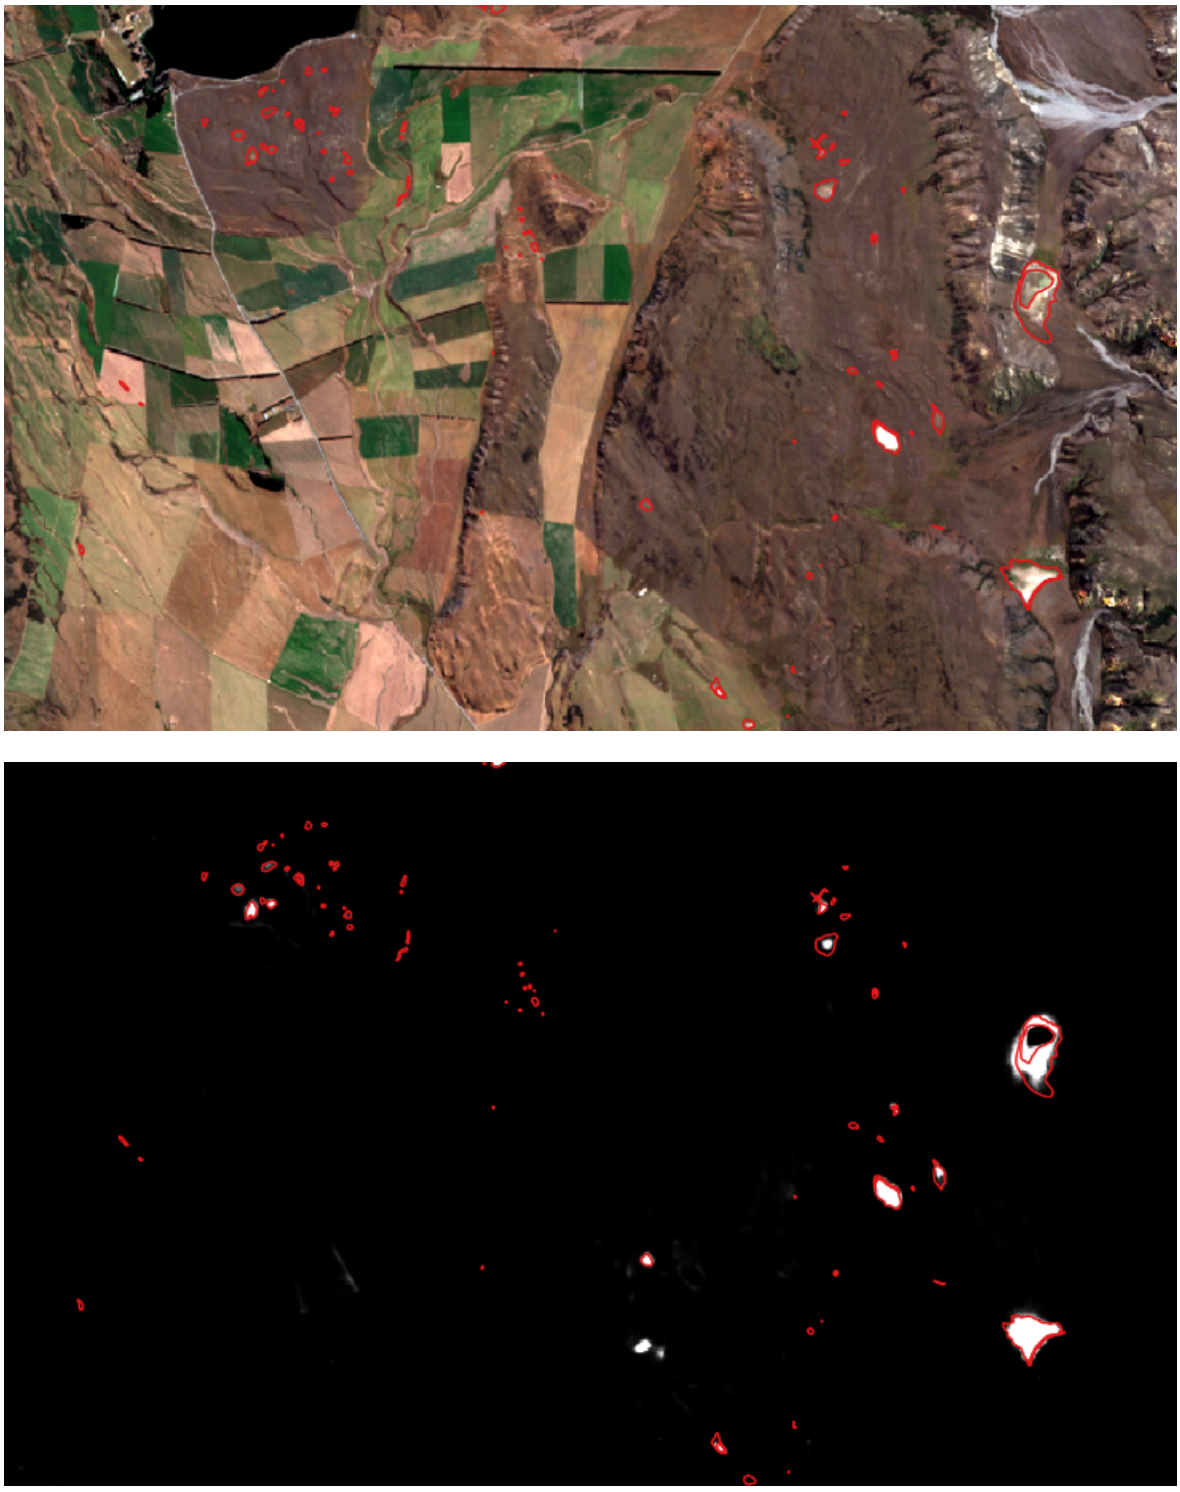

Figure 7. Sentinel-2 image (top) and model output (bottom) for the test area used by McMillan and Wiser (2019).

We note that the available training data from McMillan and Wiser only covers a small area, which limits the likely generalisability of the model. However, the model is clearly detecting at least the larger of the mapped wetlands.

We also tested the ability of the model to generalise to areas outside the training areas by training on two of the three study areas and then observing how well it detects wetlands in the third area. Figure 8 shows the prediction for the area south of Lake Heron when none of the imagery from this area was used for training.

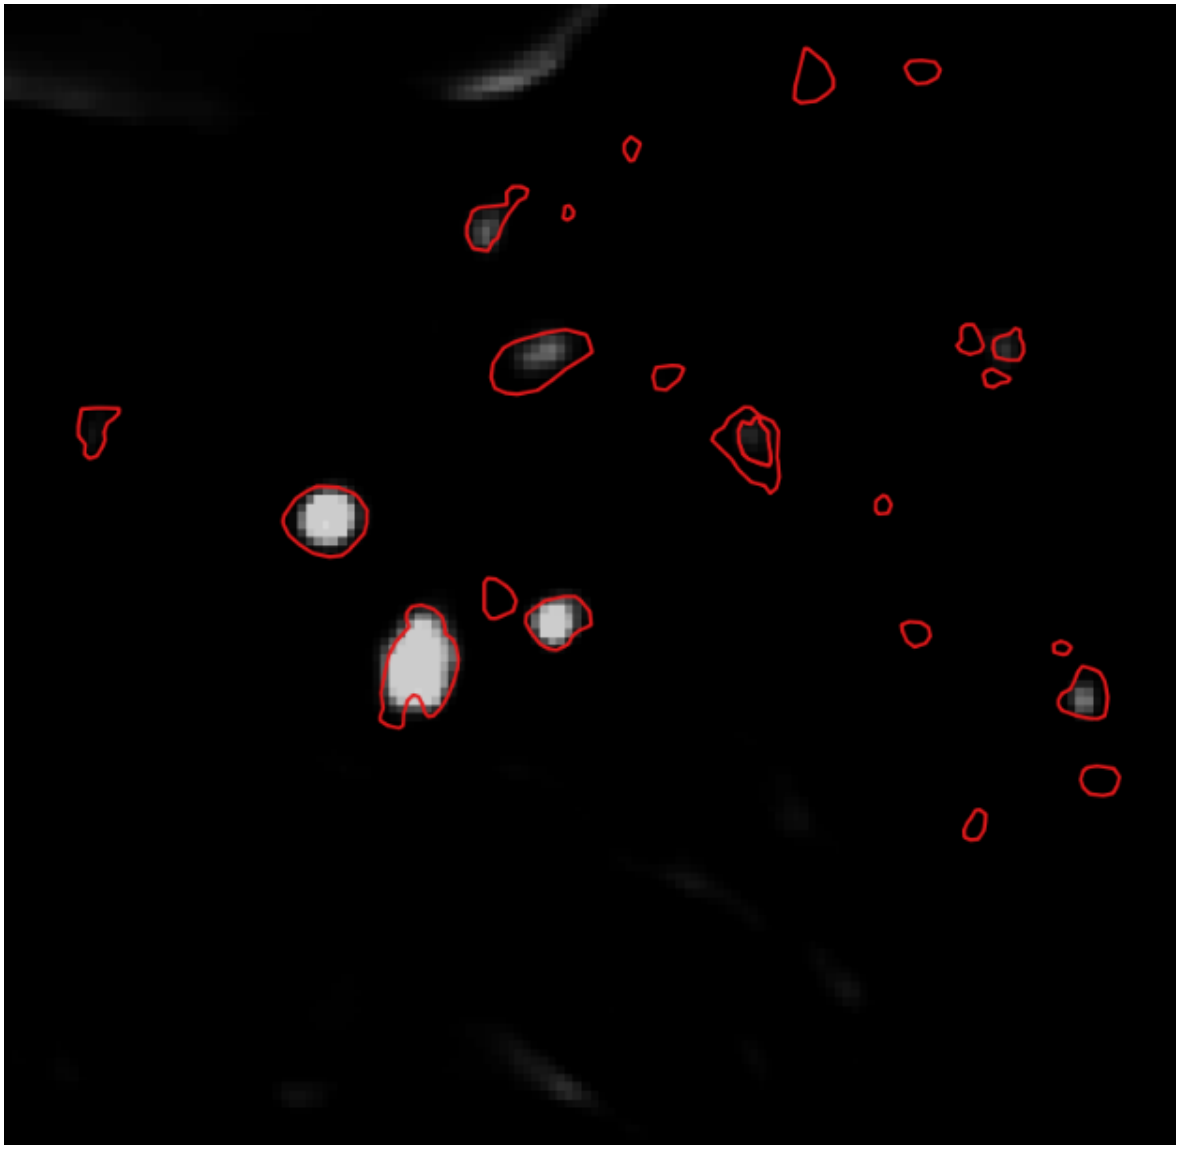

Figure 8. Model output for area south of Lake Heron.

Incorporation of monthly wetness layers

Since the predictions from the March and September image pairs were an improvement over the single image but still contained substantial numbers of false positives, we tested whether adding additional wetness information improves the prediction. This was done by including the 12-monthly wetness layers, as described in the previous section. These were stacked with the March and September Sentinel-2 10-band images to give a 32-band image stack, which was then used to train and test a new U-Net model.

Detecting kettle-hole ephemeral wetlands in the CSIGG

The original aim of McMillan and Wiser (2019) was to map kettle-hole ephemeral wetlands in the central South Island glacial geomorphology (CSIGG) area studied in Barrell et al. 2013. Figure 9 shows the area.

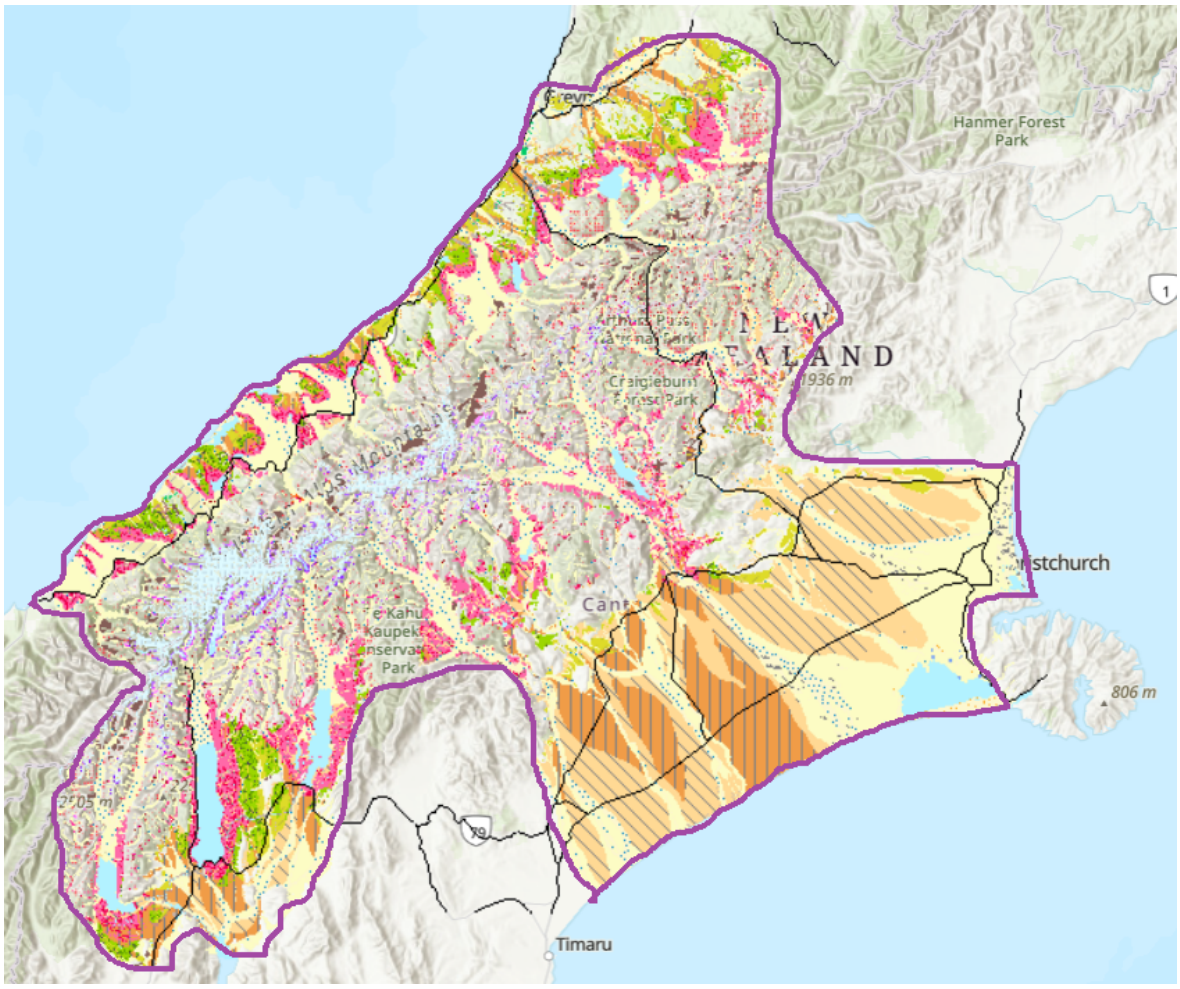

Figure 9. Central South Island glacial geomorphology area, as mapped by Barrell et al. (2013).

We used the U-Net models that produced the most promising outputs to generate probability rasters for the entire CSIGG.

McMillan and Wiser (2019) list the coordinates of known ephemeral wetlands in the CSIGG (mostly kettle holes) obtained from the literature; this was used as an initial cross-check to visually assess the method's ability to detect ephemeral wetlands from outside the three training areas. We also inspected a sample of areas where potential wetlands were detected but are not currently mapped and assessed their veracity.

4.3.2 Detecting pastoral wetlands using deep learning

Wetlands in pasture areas are difficult to identify from remote imagery. However, many pastoral wetlands also undergo seasonal changes that can assist their discovery. We tested whether the same multi-image approach used for ephemeral wetlands might be used to detect pastoral wetlands, using the Waikato region as a test case.

We had access to an excellent training data set for this. Wetlands in the Whangape catchment were extensively mapped in 2023, down to a size of 0.05 ha, which is the general minimum size for mapping wetlands under the NPS-FM. We used this area to train a U-Net model, again using imagery from the driest (February) and wettest (October) months. We also experimented with adding the 12-monthly wetness layers produced in the process-based wetland mapping component of this work.

We tested the model on images from the Whangape catchment that were excluded from the training set, and then generated and assessed a prediction for the entire Waikato region. Note that this larger prediction includes many areas that are not pastoral and are therefore expected to generate false positives. Also, the wetlands already mapped across the Waikato are of many different types (including intact peat bog, etc), so we would not necessarily expect them all to be detected by the model. We also assessed the model on a sample of candidate pastoral wetlands from the Auckland region.

5 Results

5.1 Wetness-based detection of potential ephemeral wetlands

5.1.1 Time series behaviour of ephemeral wetlands

Figure 10 depicts the NDVI time series for two test areas: one over an ephemeral kettle-hole wetland and another that is not classified as ephemeral wetland. Overall, the index shows similar timings in highs and lows, but with more pronounced extremes for the kettle-hole wetland, which could mean that it seasonally filled up with water (very low NDVI) in winter but is covered in more vegetation during dry summer (high NDVI).

We contrasted dozens of such examples, but the absolute values vary substantially between wetlands, and the timings are different (i.e. the time of the year when the NDVI increases or drops). It was thus not possible to classify ephemeral wetlands using a spectral index alone. This also meant that annual median images of Sentinel-2 are not useful for identifying wetlands; instead, monthly averages or seasonal maps are needed.

We also investigated the potential use of Sentinel-1 C-band radar images, but due to the inherent radar speckle, individual images were too noisy at these fine scales (<0.1 ha).

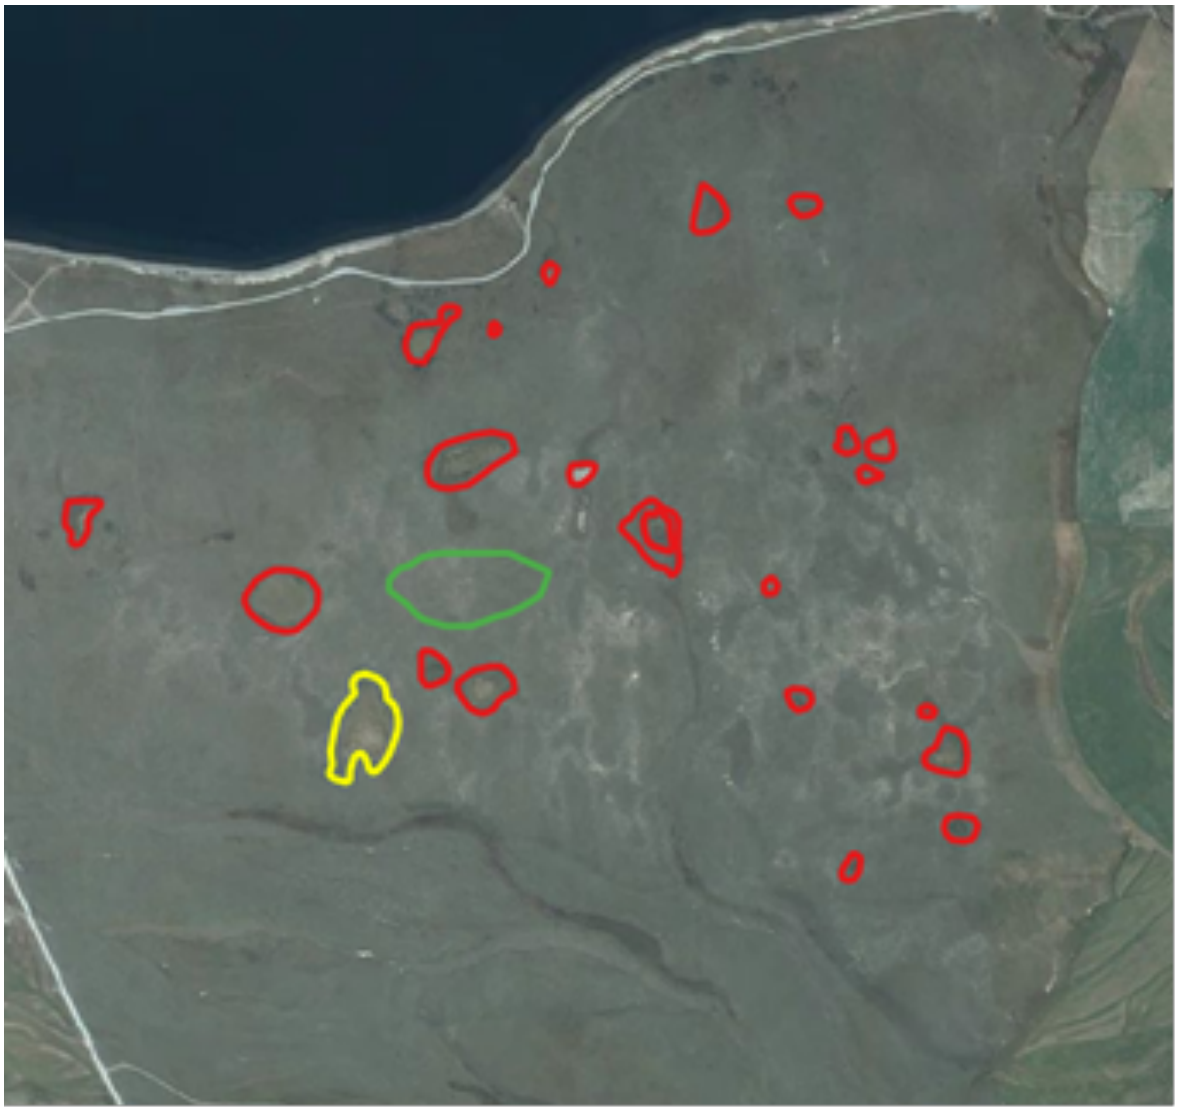

(a)

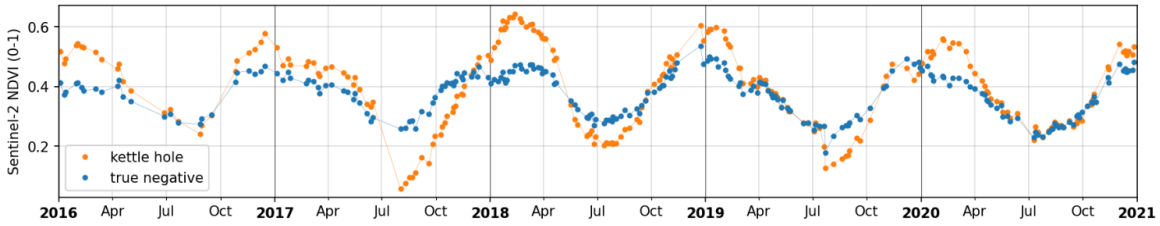

(b)

Figure 10. (a) Example of a kettle-hole ephemeral wetland from the DOC ground-truth data (yellow polygon) and a non-kettle-hole polygon area (green polygon) next to it. (b) The Sentinel-2 NDVI time series shows strong seasonal behaviour in both time series. The extremes in the NDVI are more pronounced for the kettle hole (orange series), with minimum close to 0.1 and maximum over 0.6, compared to the non-wetland polygon (blue series).

5.1.2 National water frequency mapping

The water classification model was applied to each individual HQ Sentinel-2 image for the year 2021. The number of water observations ('water counts') at the pixel level was then computed monthly and annually, which gives a spatial representation of water and inundation frequency for all New Zealand (Figure 11). These layers represent the most comprehensive and spatially detailed surface water map created to date, and they map a variety of water bodies, including different-coloured lakes, braided rivers, ephemeral ponds, and irrigation basins (Figure 12).

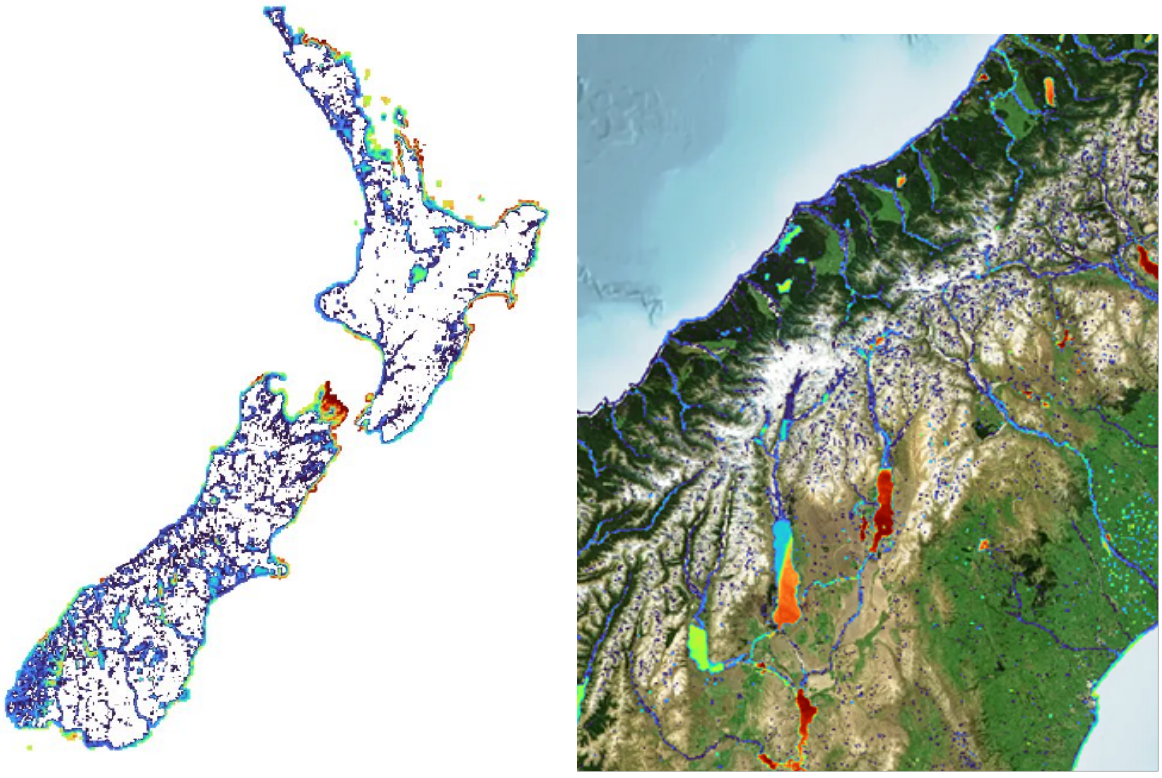

Figure 11. Water frequency layer for New Zealand in 2021. The maps show the water counts as observed by the Sentinel-2 sensor. Light green and blue indicate fewer water observations, and orange and dark red more water counts.

Notes: Due to overlapping orbits, cloud coverage, and cast shadow (from mountains), the absolute number varies for different parts of the country. This uneven pattern is visible in Lake Pukaki.

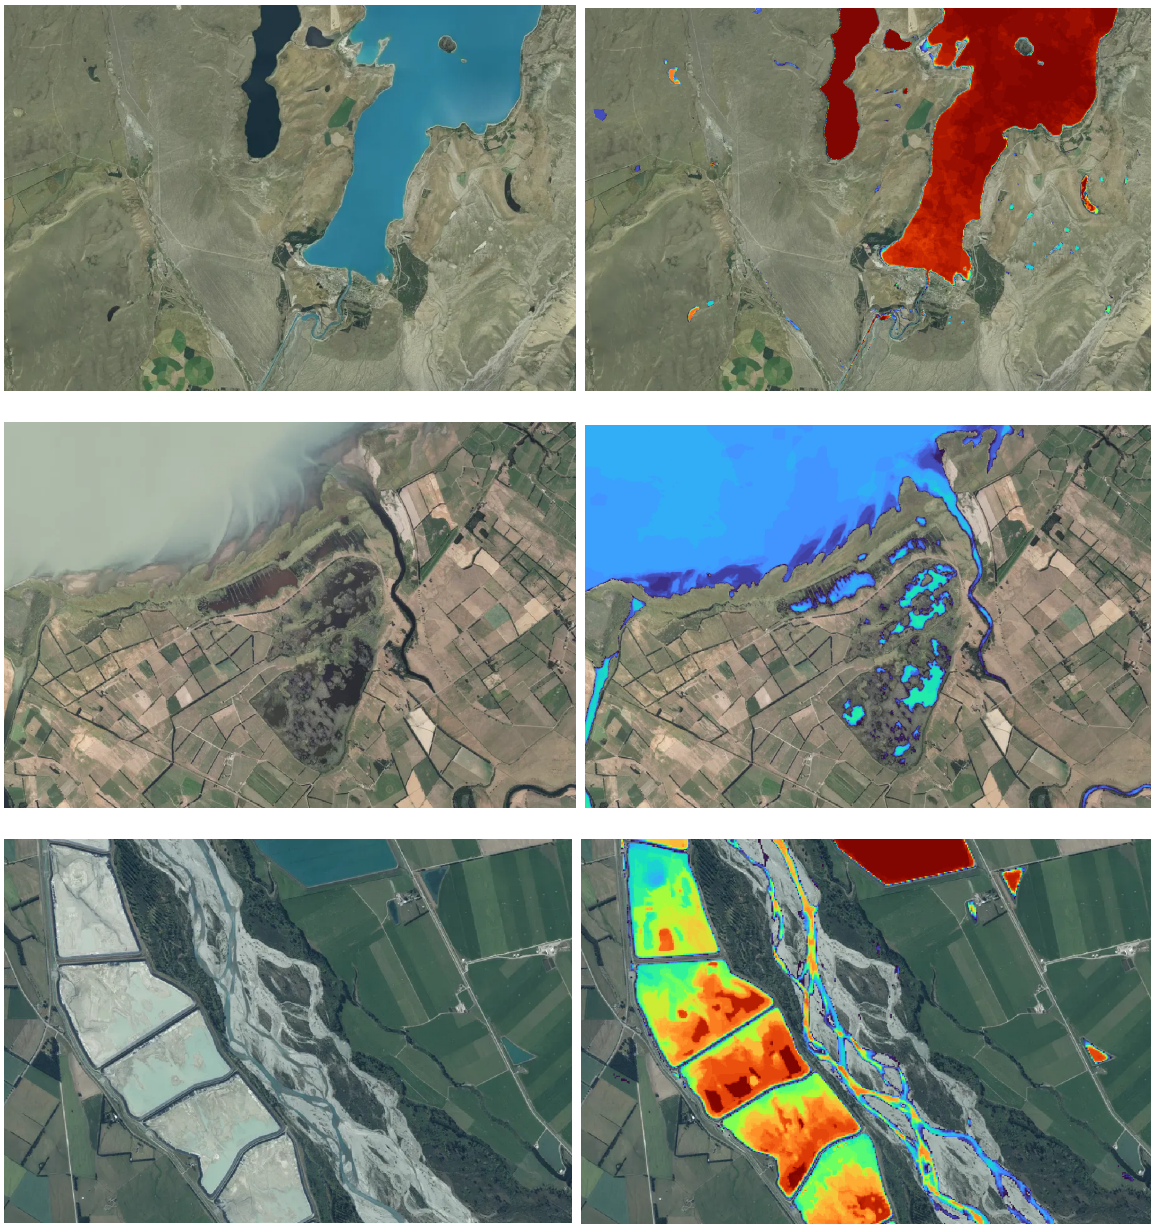

Figure 12. Aerial imagery (left) and overlaid water frequency (right) for different types of water bodies. Top panels show permanent lakes Tekapo and Alexandrina (red), as well as seasonal ponds (lighter colours). Middle panels shows the Wairio wetland in Wairarapa with many seasonally inundated ponds and main flow patterns of the larger lake Wairarapa. Bottom panels depict a frequently changing braided Rangitata River and water storage basins, which are part of the Rangitata South Irrigation Scheme (RSIS).

Detecting ephemeral wetlands

The main objective of this study is to make use of Sentinel-2 water observations to help detect ephemeral wetlands. The previous section demonstrated the ability of the frequency layer to map permanent, semi-permanent, and short-term water features in different settings. Small, ephemeral wetlands are, by nature, hard to spot and delineate. Aerial photography and satellite imagery must usually be taken at the right time of the year or season to capture ponding behaviour. In contrast, the Sentinel-2-derived annual water frequency layer located 210 out of the 506 confirmed ephemeral wetlands from the MWLR/DOC ground-truth data set (41%) without special training or optimisation (Figure 13).

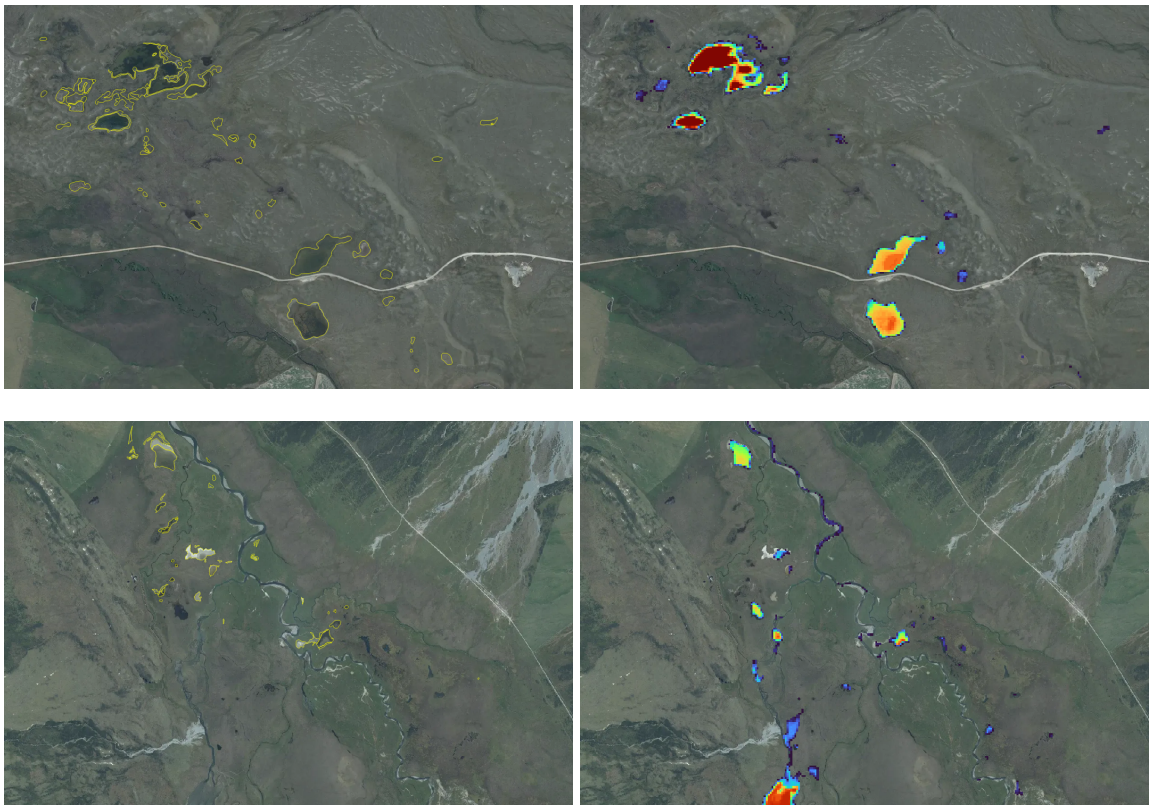

Figure 13. Potential use of the water frequency layer to delineate ephemeral wetlands. Left-hand panels show aerial imagery with some ephemeral wetlands in Canterbury currently inundated. Right-hand panels show them identified by water frequency results, demonstrating that water frequency is a useful metric to locate and estimate the extent of these seasonal wetlands.

Comparing the Sentinel-2 water frequency to aerial photography is not appropriate, however, because the photography captures spatial detail at 30 cm pixel resolution whereas Sentinel-2 operates at 10 m. The frequency layer is able to map individual clear water pixels but still struggles to identify small, vegetated ponds due to spectral mixture effects from adjacent pixels. Also, some of the manually identified ephemeral wetlands from the ground-truth data sets are not recognised, even though they appear large enough. This may be due to the limited time series of the frequency layer, which only comprised image dates from the year 2021. Also, most of the MWLR/DOC polygons were drawn in earlier years dating back to 2015 (see McMillan & Wiser 2019), so some may be gone or were not inundated in 2021.

Figure 14 demonstrates that the frequency layer successfully detects the ground-truth polygons, even in situations where the aerial photograph does not show any indication of wetlands.

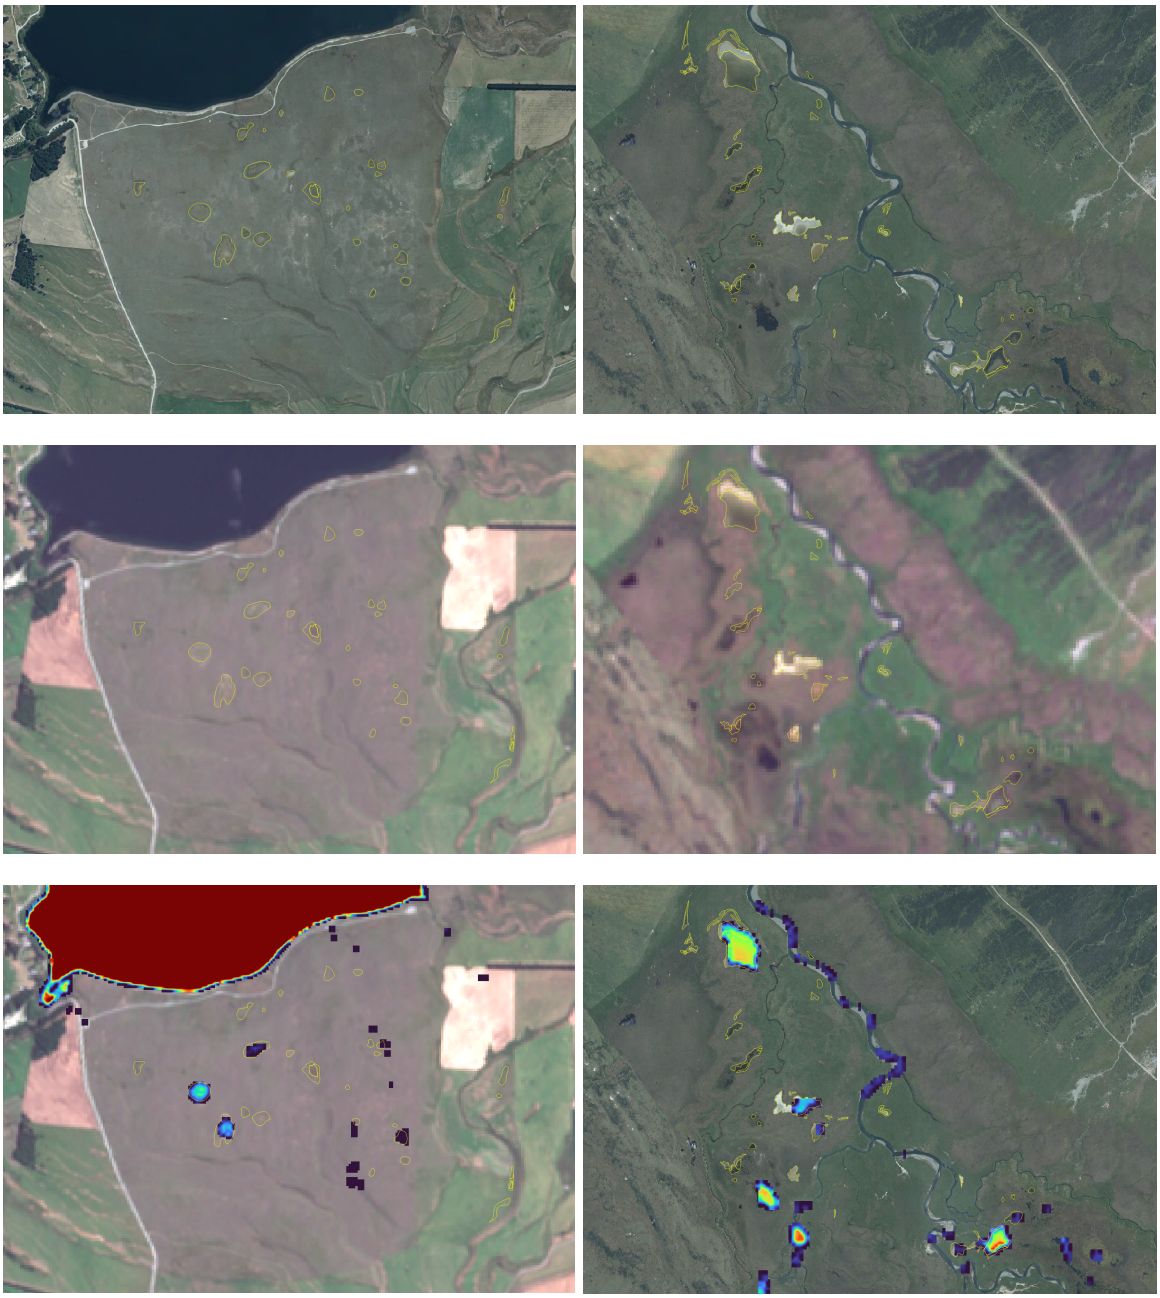

Figure 14. Potential use of the water frequency layer to delineate ephemeral wetlands in imagery of different spatial resolutions and from different observation platforms (aerial and satellite). Top panels show the high-resolution aerial image and ground-truth polygons. Middle panels show the same areas with Sentinel-2 imagery. Bottom panels show the Sentinel-2 water frequency on the aerial photographs, with many wetlands correctly identified.

Targeted training for small, vegetated ponds

The Sentinel-2-derived water frequency layer tends to struggle with small, vegetated ponds. The classification model is therefore often unable to detect such small ponds in certain areas. We tried adding targeted training data (i.e. point locations over small ponds) and re-trained the logistic regression model. The new model located small ponds more reliably, but was also more prone to false predictions over forested areas with similar spectral signatures (Figure 15).

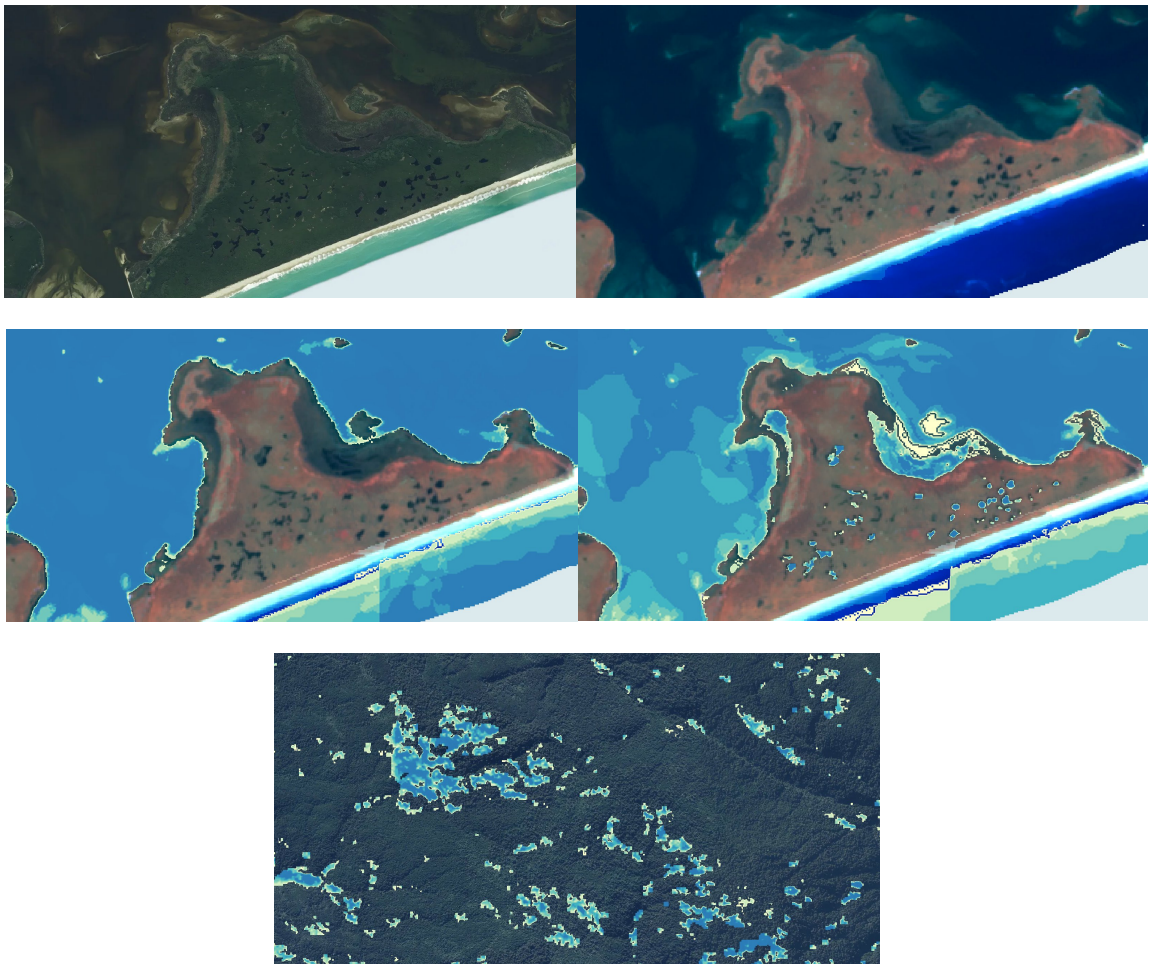

Figure 15. Aerial photo (top left) and Sentinel-2 image (top right). The original water classifier is unable to detect small, vegetated ponds (middle panel, left). The updated model with targeted training data performs better (middle panel, right), but also introduces more false positive predictions over forests with similar spectral signatures such as the detected areas in the bottom panel.

5.1.3 Water body detection in high-resolution aerial photography

Model performance

The U-Net model for detecting water from aerial imagery alone was trained for only 1 hour and 21 minutes but quickly learned to classify water well. The training accuracy (F1-score) of water at the highest point was 0.984, together with a validation of 0.961, which is extremely high. We stopped the model training after 175 epochs, because no further improvement could be observed and overfitting to the training data became stronger (Figure 16).

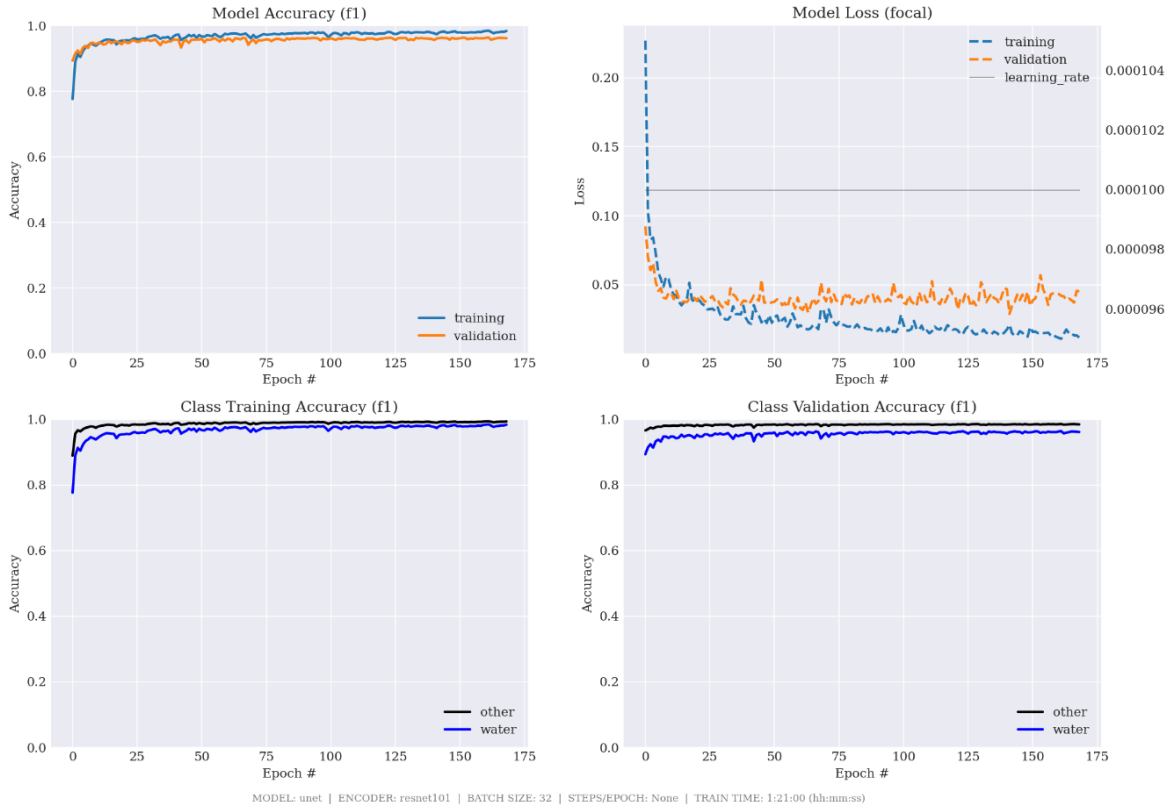

Figure 16. U-Net model training curves showing the accuracy and loss for water classification in aerial photography over many epochs. The accuracy is very high after 25 epochs and then keeps slowly improving until epoch 75, after which no further improvement can be seen.

Prediction quality

We applied the water detection model to the entire Greater Wellington and Hawke's Bay image data sets and created two corresponding regional water masks (Figure 17). Overall, this captures all types of water bodies well, and the spatial detail is much higher than the Sentinel-2-based water layer (detailed view Figure 18).

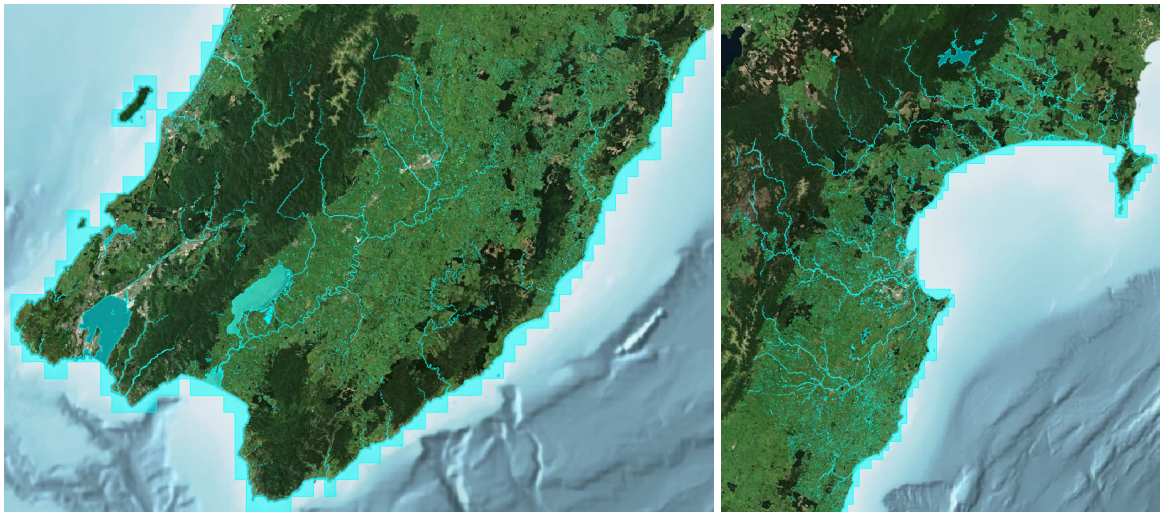

Figure 17. Deep learning-based water mask for Greater Wellington and Hawke's Bay Regions.

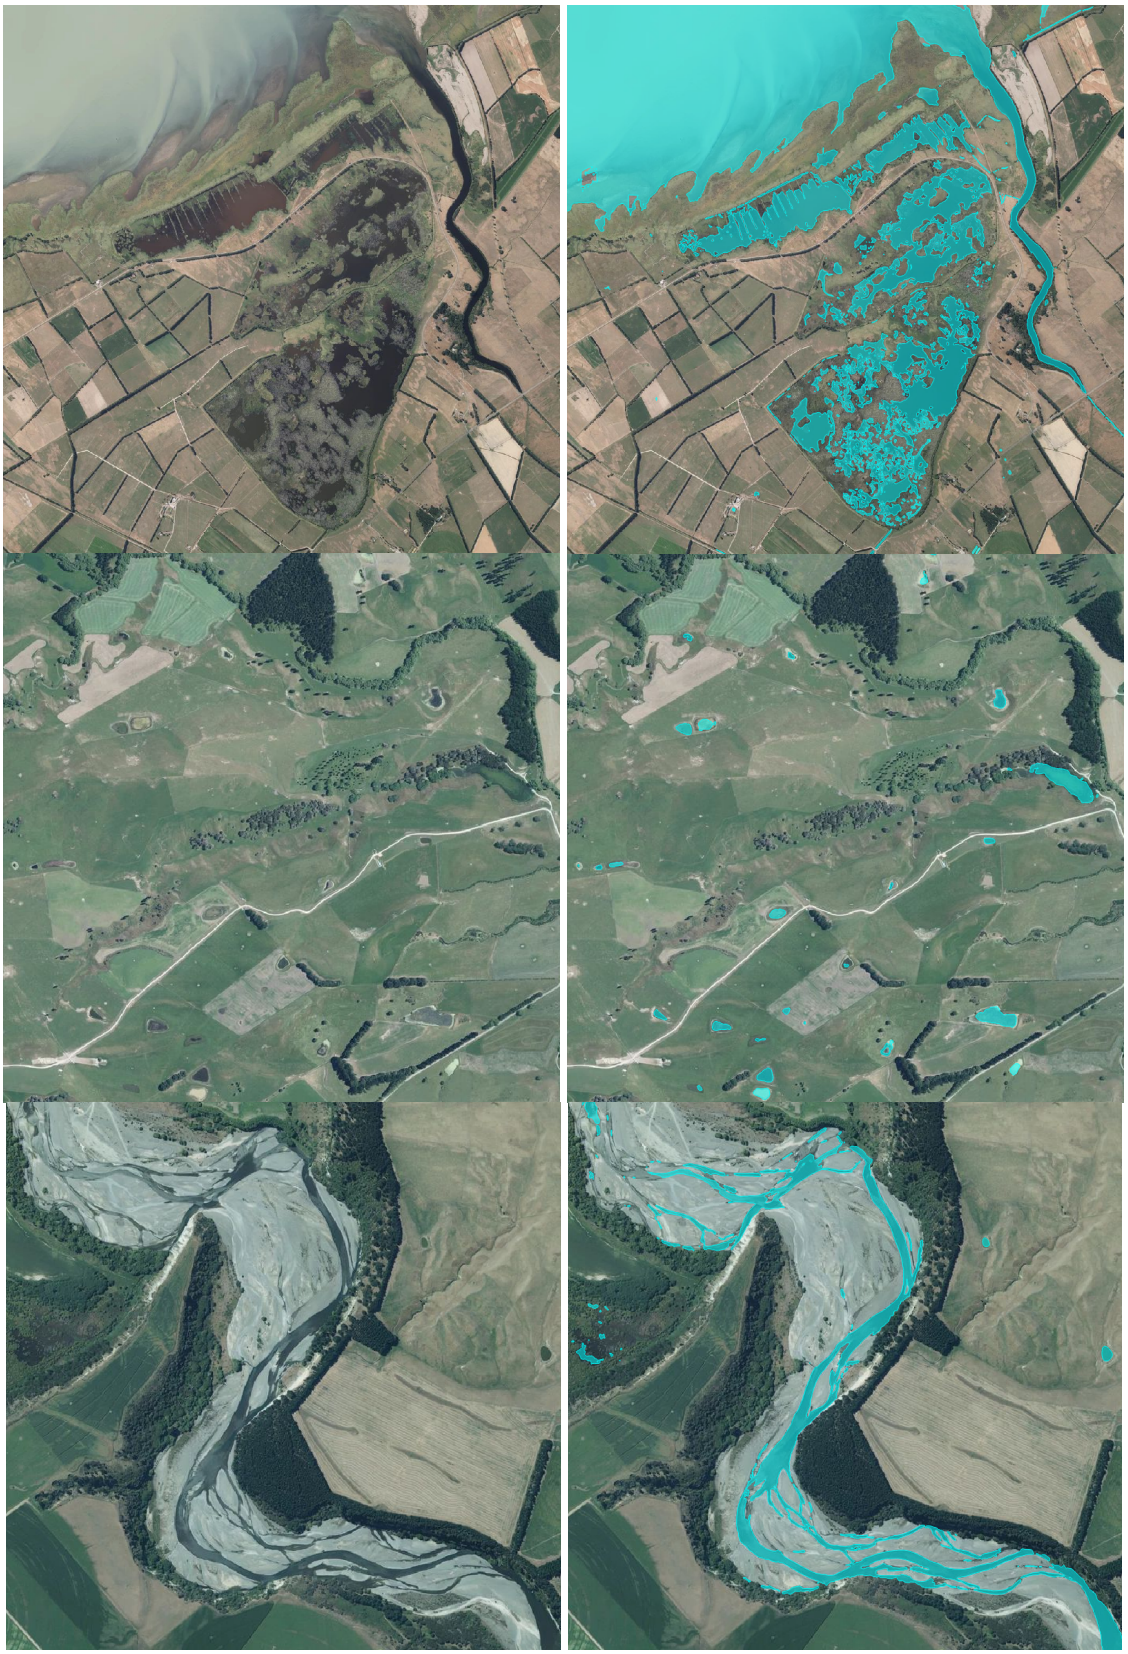

Figure 18. Examples of the deep learning-based water mask for Greater Wellington and Hawke's Bay regions showcasing the spatial detail. Top panels depict the Wairio wetland in Wairarapa, middle panels show kettle holes and small ponds (locality), and bottom panels show complex river systems that are part of the Ngaruroro River in central Hawke's Bay.

5.2 Image-based detection of potential wetlands

5.2.1 Detecting kettle-hole ephemeral wetlands using deep learning

Detecting known kettle-hole ephemeral wetlands

We trained two models and used them to generate a probability raster for each of the three study areas. The first model was trained using just the March and September 2020 Sentinel-2 imagery ('two-image model'), while the other included the 12 monthly wetness rasters ('combined model'). We visually compared the outputs from the two models to the kettle-hole wetlands mapped by DOC and MWLR to assess sensitivity and precision.

Figure 19 to Figure 21 show outputs for the two-image and combined models for sample test images from each study area. In all three cases the combined model has lower sensitivity but significantly higher precision (i.e. fewer false positives and more accurate wetland shapes). Without an exhaustive map it is not possible to definitively select one model over the other; in practice this may depend on how the output raster is used. Both models demonstrate that they detect kettle-hole wetlands in new imagery to a reasonable degree, at least for imagery taken close in time to the training images.

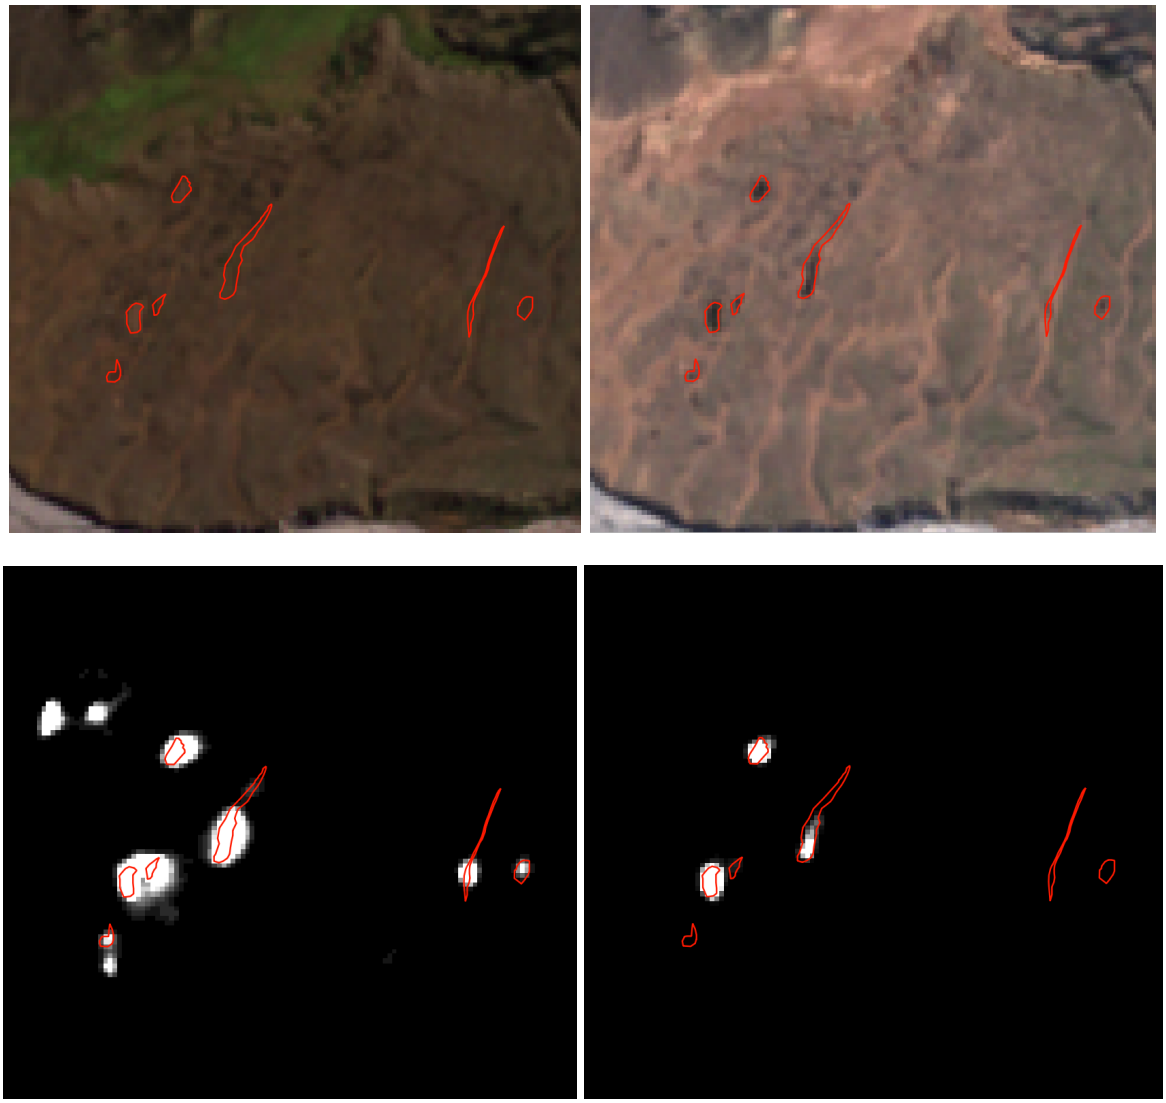

Figure 19. Model outputs for a test image in study area 1 (NZTM 1454443, 5190941). Top: 2020 Sentinel-2 images for March (left) and September (right). Bottom: predictions for two-image model (left) and combined wetness model (right).

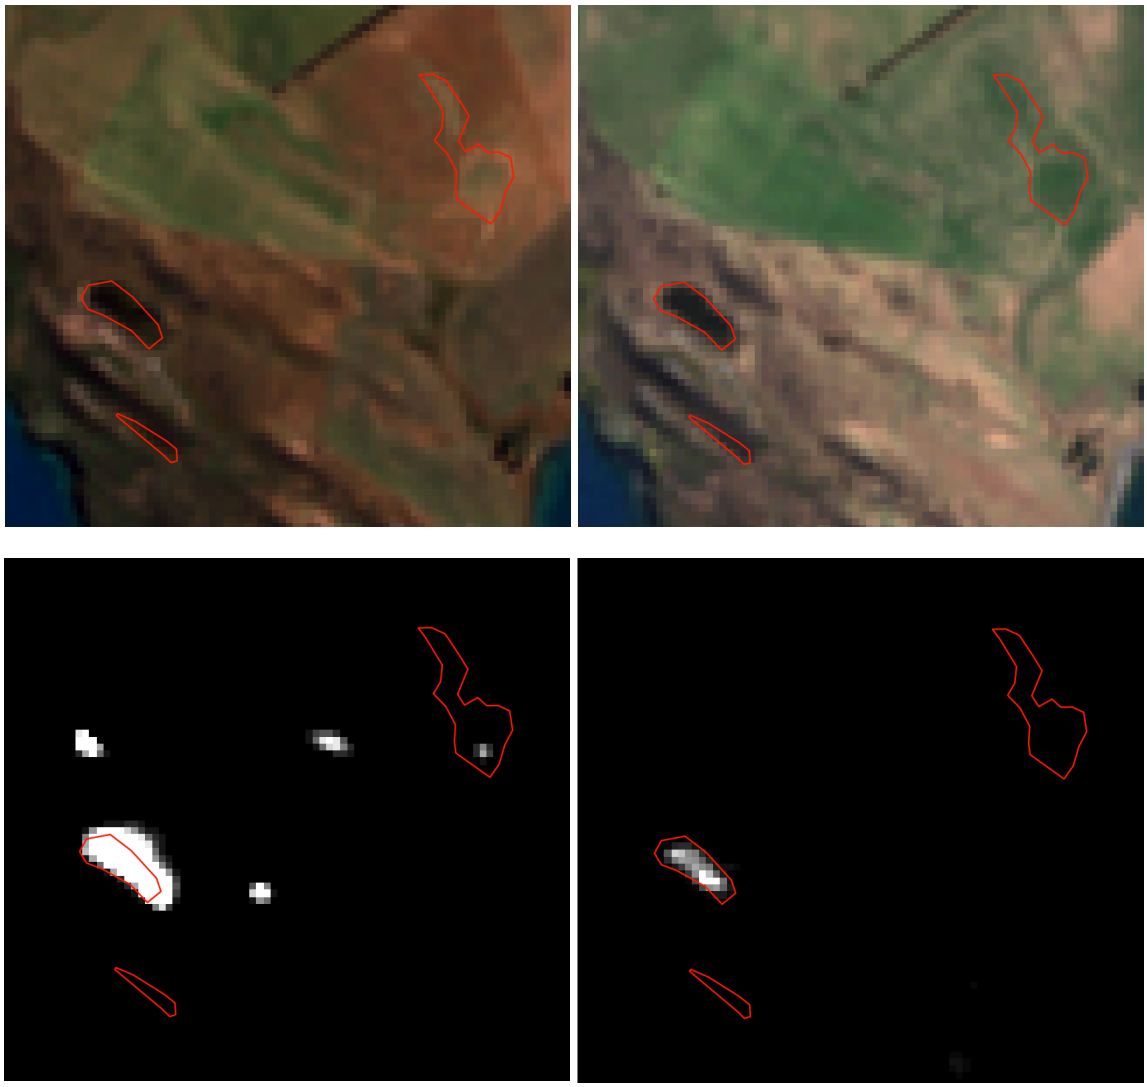

Figure 20. Model outputs for a test image in study area 2 (NZTM 1479898, 5204251). Top: 2020 Sentinel-2 images for March (left) and September (right). Bottom: predictions for two-image model (left) and combined wetness model (right).

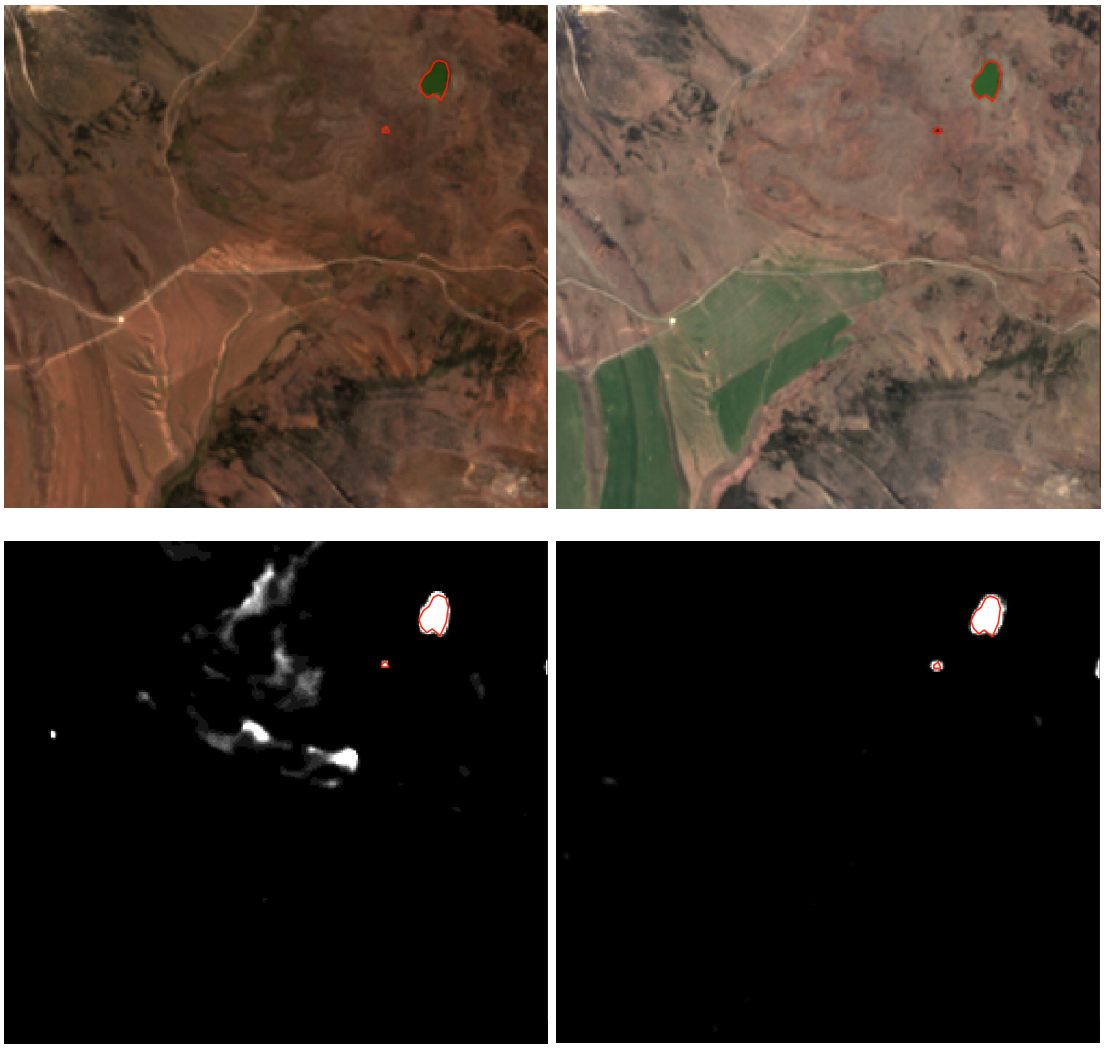

Figure 21. Model outputs for a test image in study area 3 (NZTM 1506496, 5223305). Top: 2020 Sentinel-2 images for March (left) and September (right). Bottom: predictions for two-image model (left) and combined wetness model (right).

Detecting kettle-hole ephemeral wetlands in the CSIGG

We used the two-image and combined models to generate a probability raster for the entire CSIGG. Figure 22 shows the prediction for the CSIGG using the combined model. This shows that the model has identified a significant number of potential further wetlands across the area. There is no complete data set available to test the model because the CSIGG has not been extensively mapped, so we again relied on visual inspection of sample areas to assess whether the model is likely to be useful.

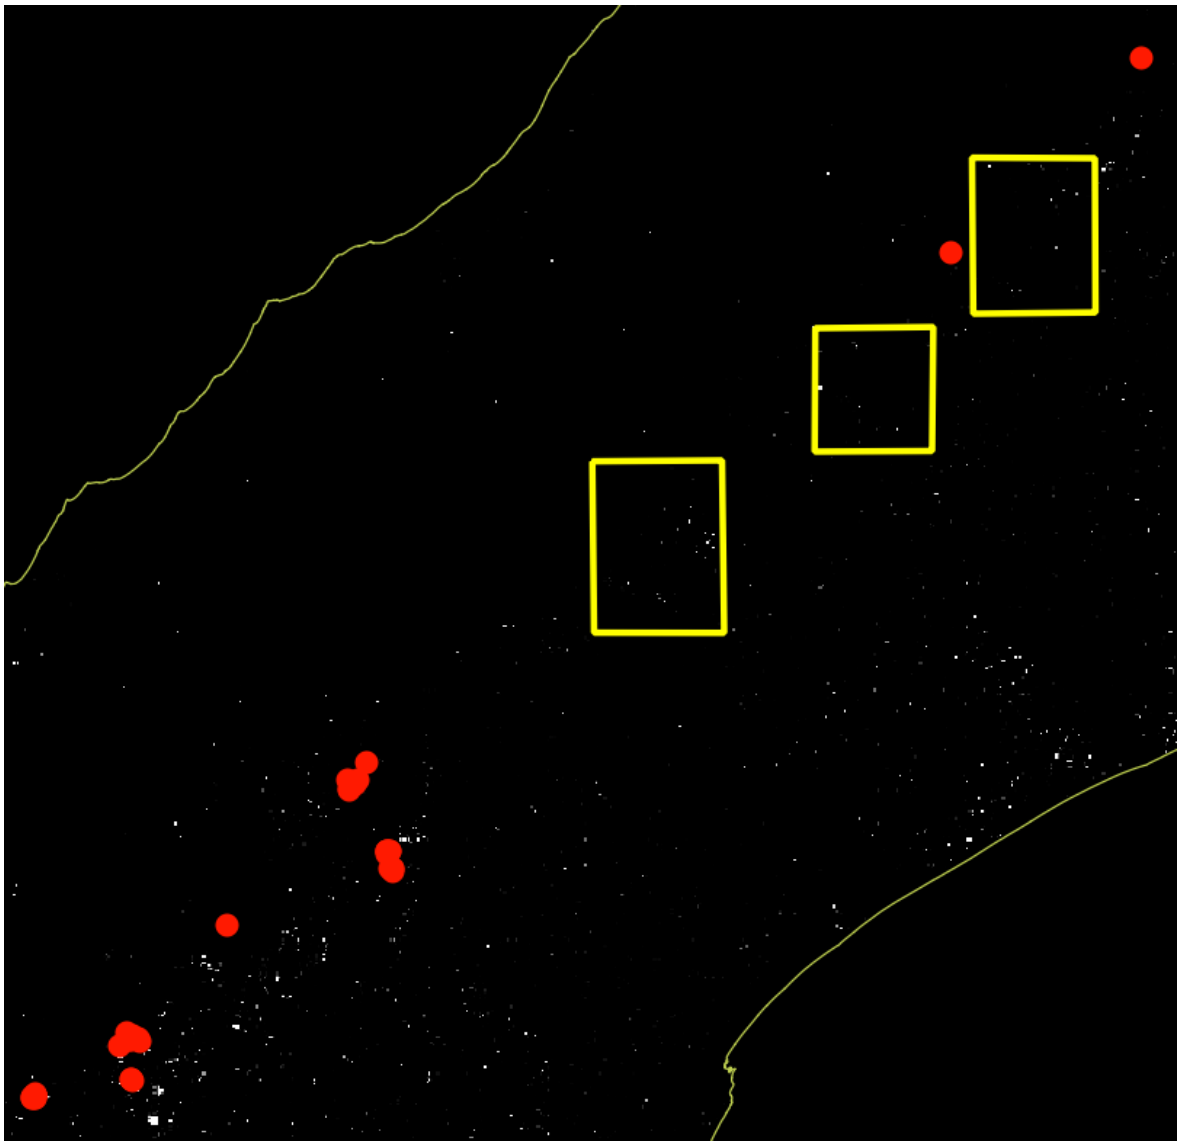

Figure 22. Combined model output for the central South Island glacial geomorphology area. Yellow rectangles indicate the areas used for training; red blobs show the locations of known ephemeral wetland areas from the literature.

To gauge the level of detection, we inspected known ephemeral wetland locations from the literature, as listed in McMillan & Wiser 2019. Figure 23 Figure 26 give some examples of the model's outputs in areas containing known wetlands. Note that black areas in some images are where no imagery was available because of cloud cover.

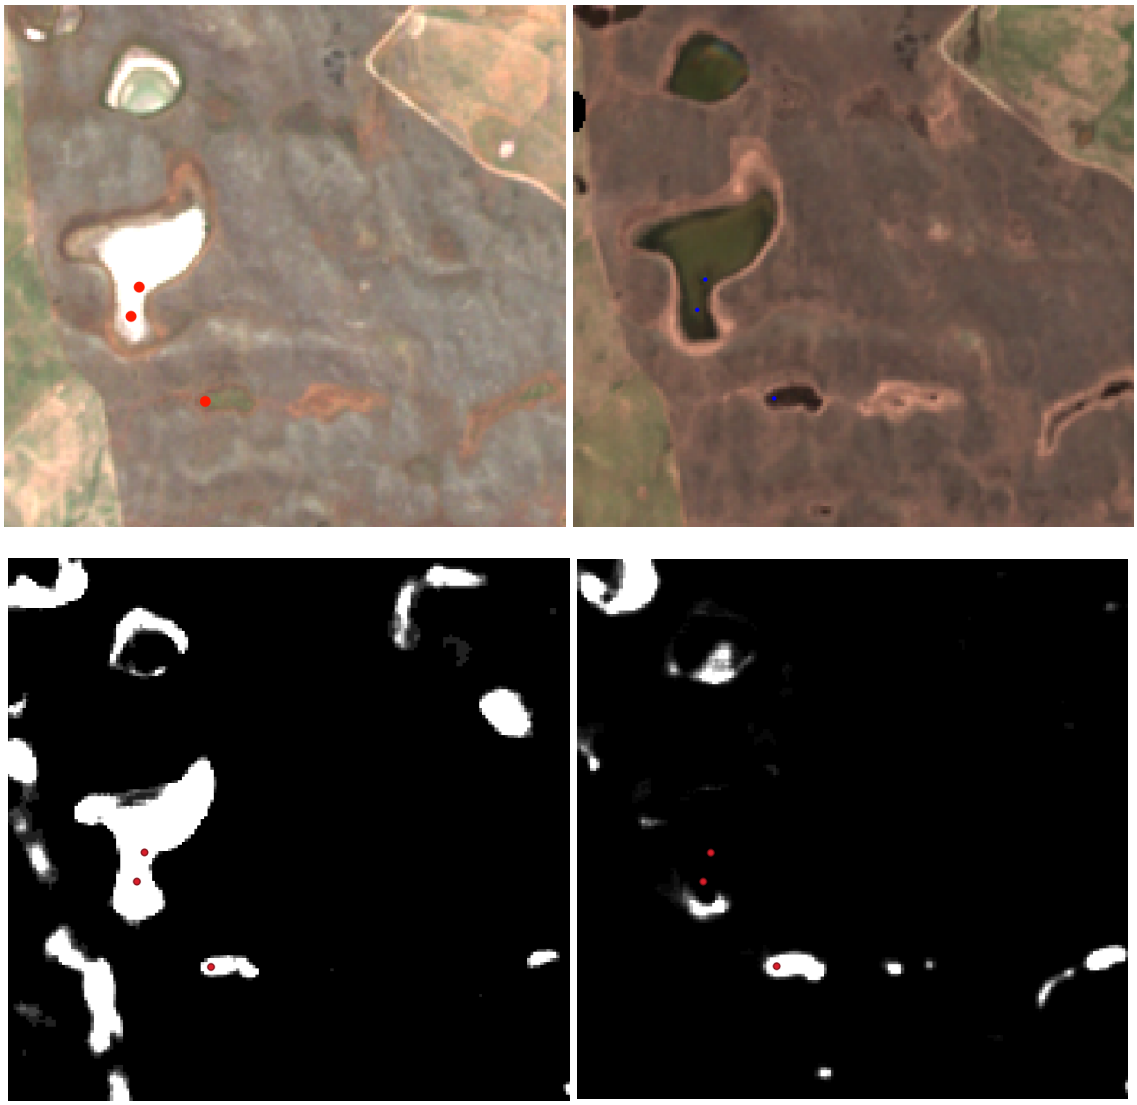

Figure 23. Wairepo kettle holes conservation area. Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

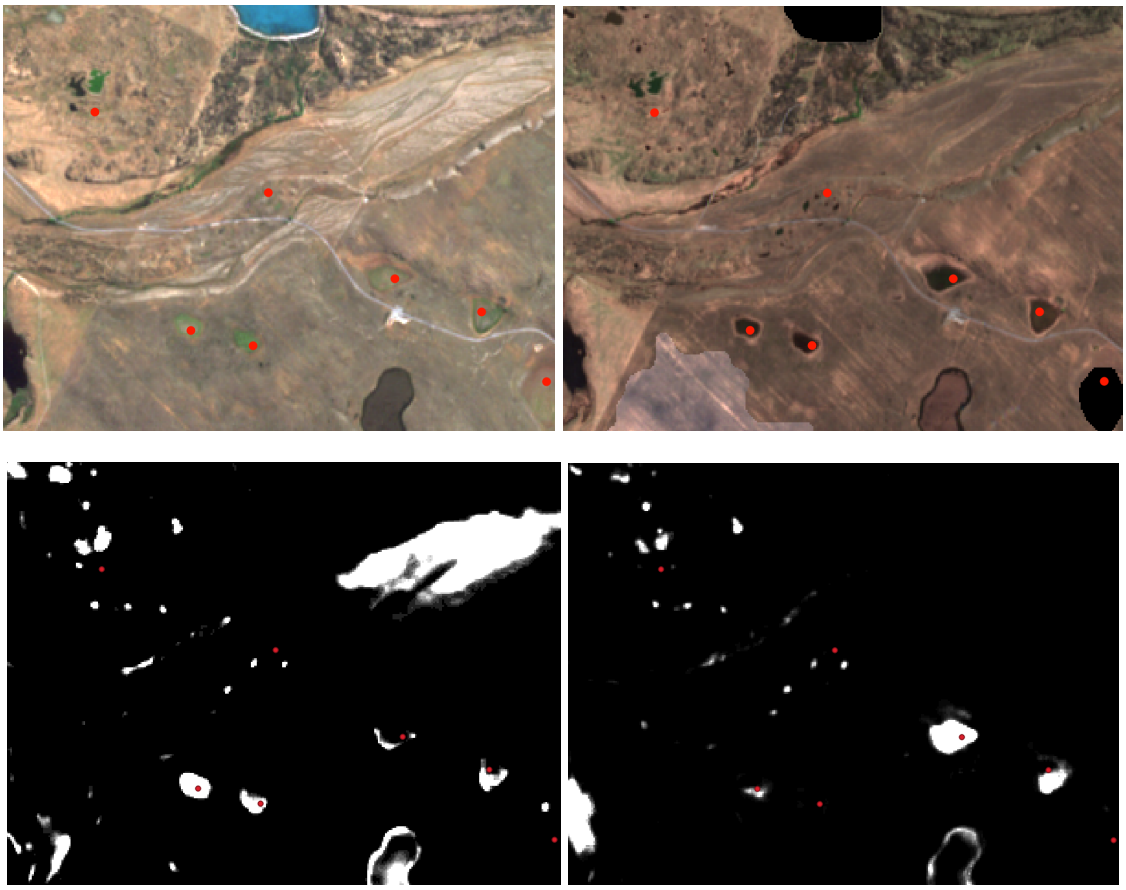

Figure 24. South of Lake Ōhau. Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

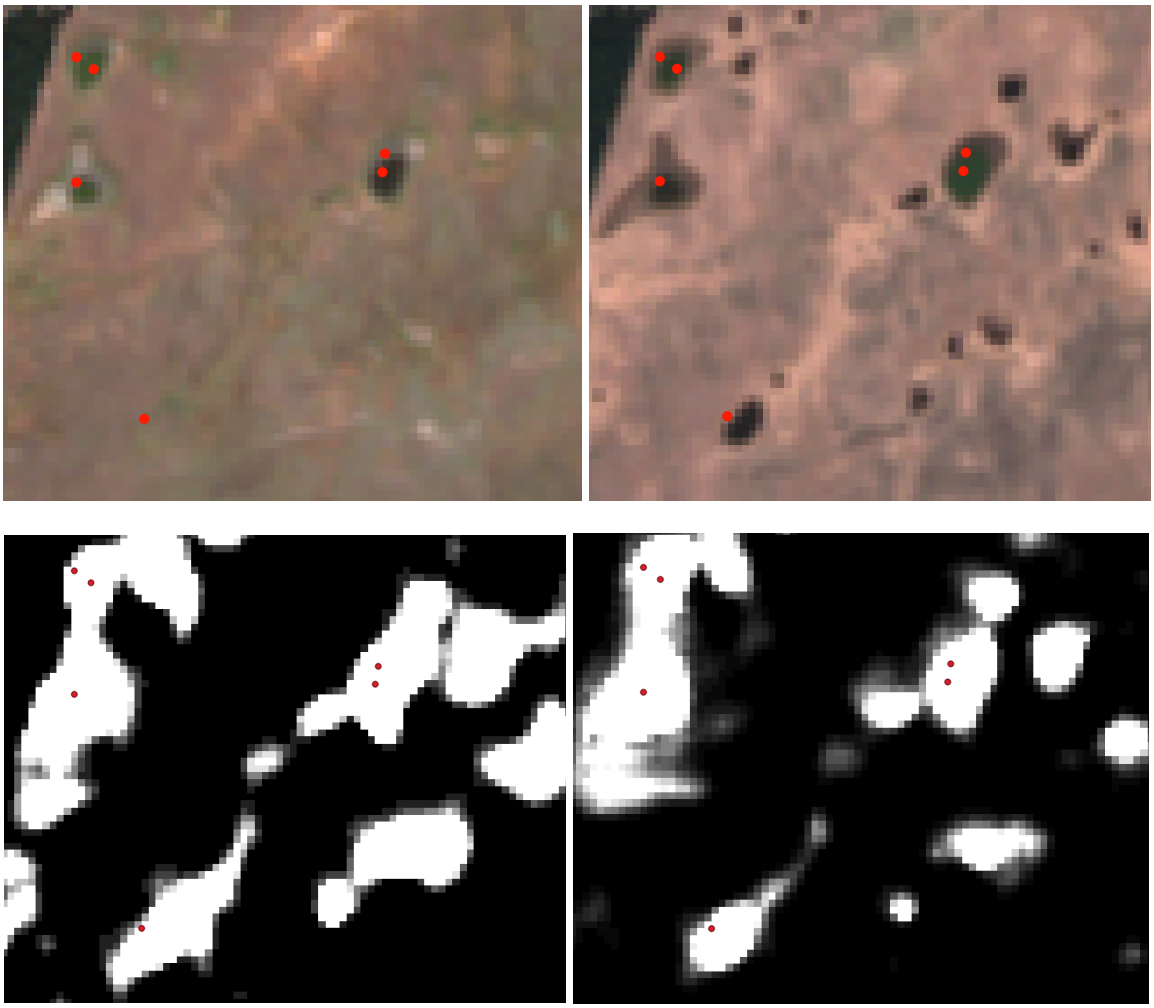

Figure 25. South of Lake Tekapo. Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for two-image model (left) and combined model (right).

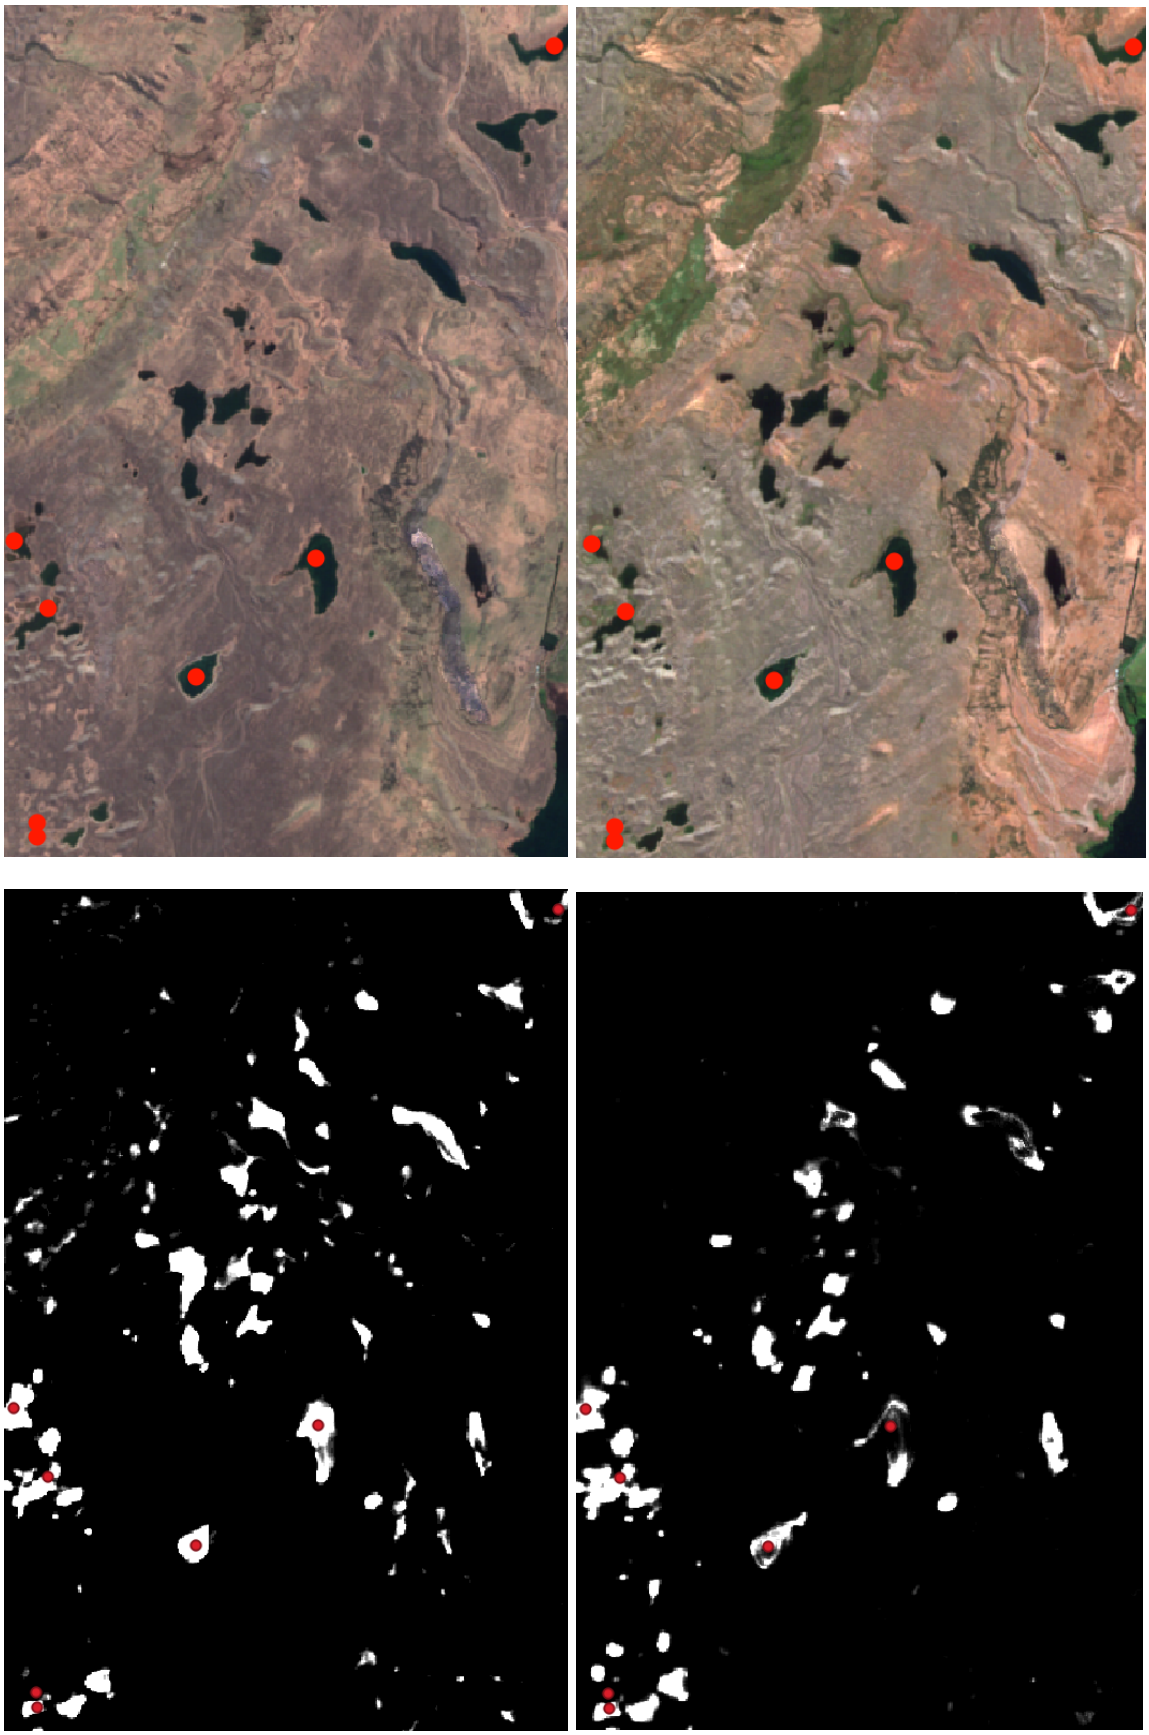

Figure 26. West of Lake Tekapo. Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

Both models show strong detection of the known wetlands. In general, the combined model appears to be more sensitive to wetland areas that undergo the most seasonal changes but misses some wetlands that remain full of water, whereas the two-image model detects these wetlands but also includes a higher level of false positives. This is expected given the strong temporal nature of the data used to train the combined model.

We also inspected areas where the models detected potential wetlands where none were previously mapped. Figures 27 to 29 show some examples where both models detected potential wetlands in the alpine areas of the CSIGG similar to, but distant from, the three training areas.

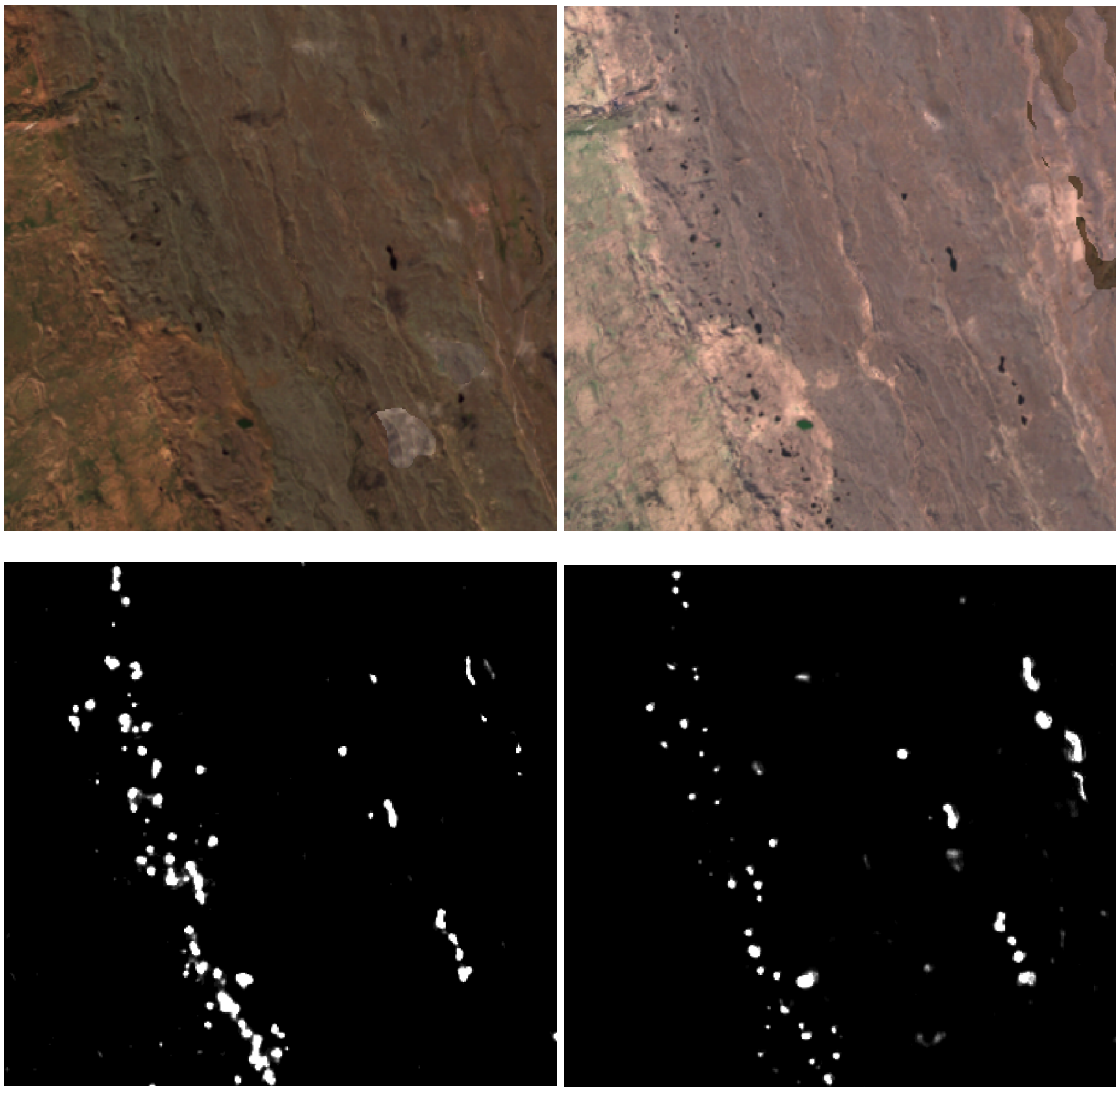

Figure 27. North-east Lake Pukaki (NZTM 1378544, 5130376). Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

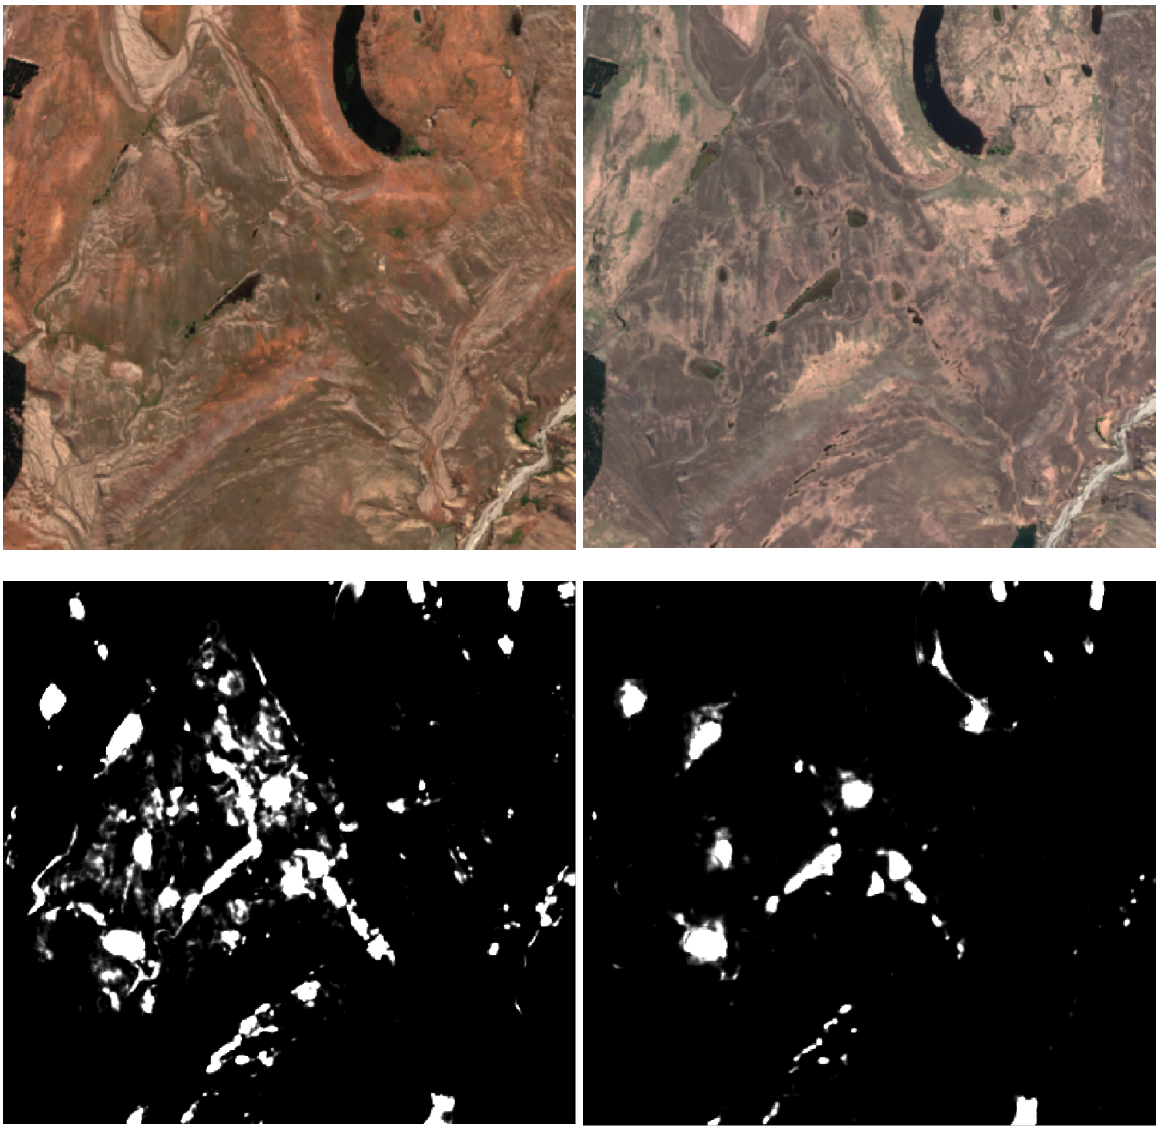

Figure 28. South-east Lake Tekapo (NZTM 1402577, 5125730). Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

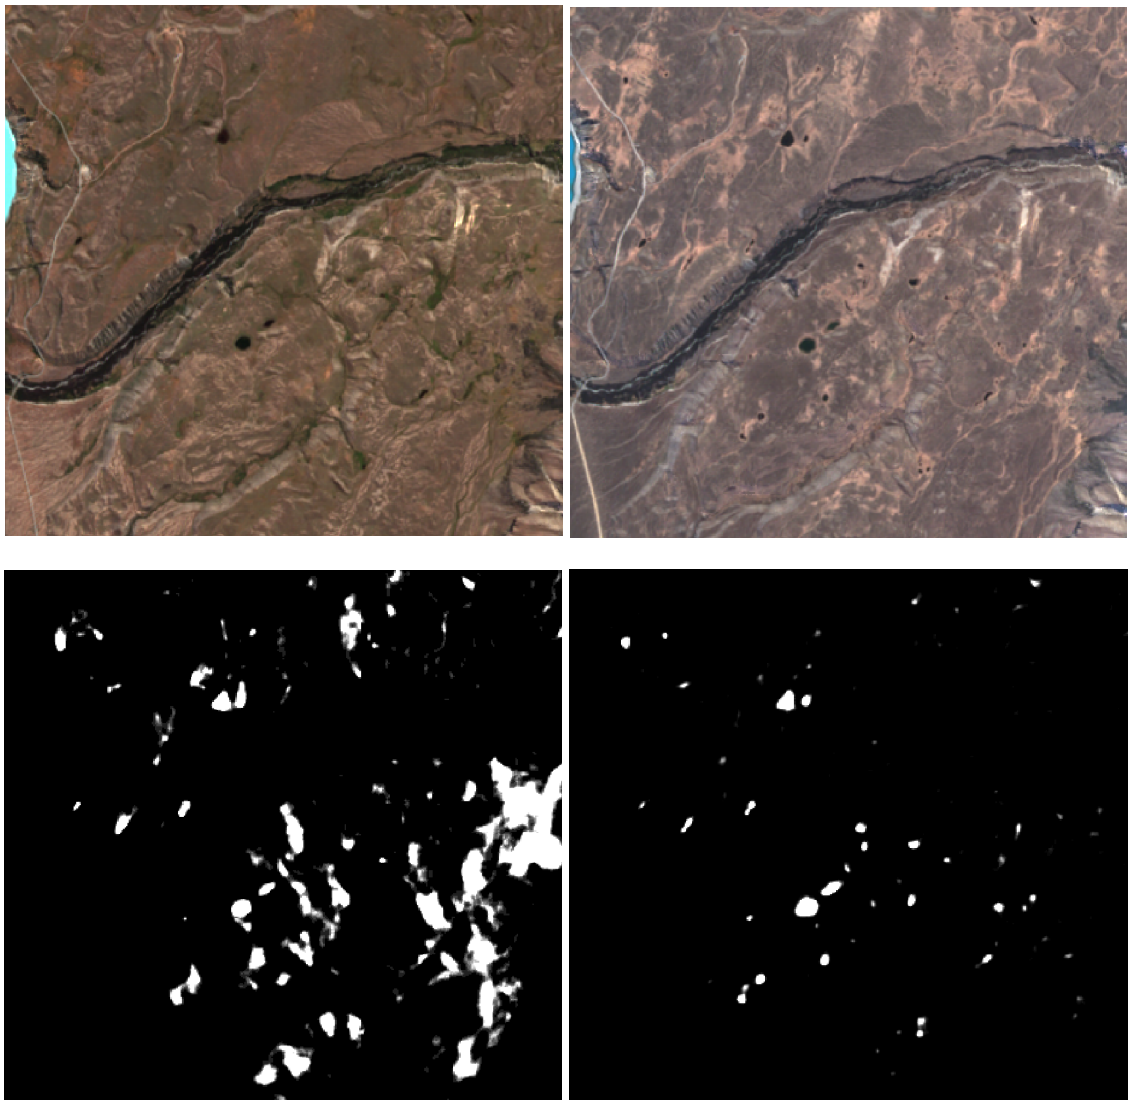

Figure 29. East of Lake Tekapo (NZTM 1405848, 5135689). Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

For alpine areas, the combined model appears to accurately capture potential kettle-hole wetlands, as well as other wetland types. The two-image model again has higher sensitivity but appears to return a significant area of false positives.

These models were trained on small samples of alpine areas. Once the terrain and vegetation change significantly, the model's accuracy falls. Figure 30Figure 32 show examples of the model results on areas with significantly different land cover, land use, and terrain.

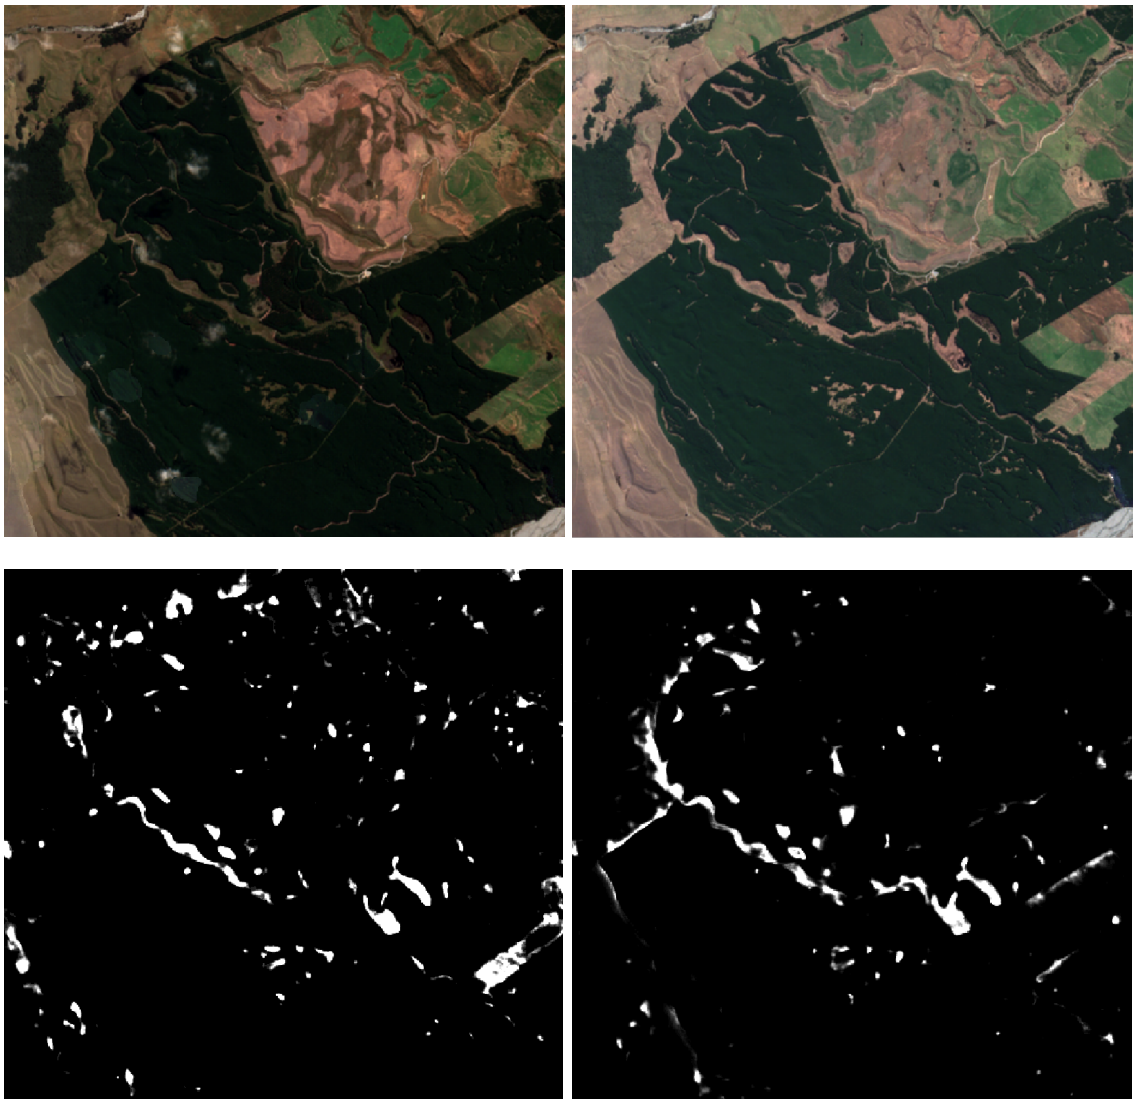

Figure 30. Forestry block between Lake Heron and Lake Tekapo (NZTM 1429024, 5160820). Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

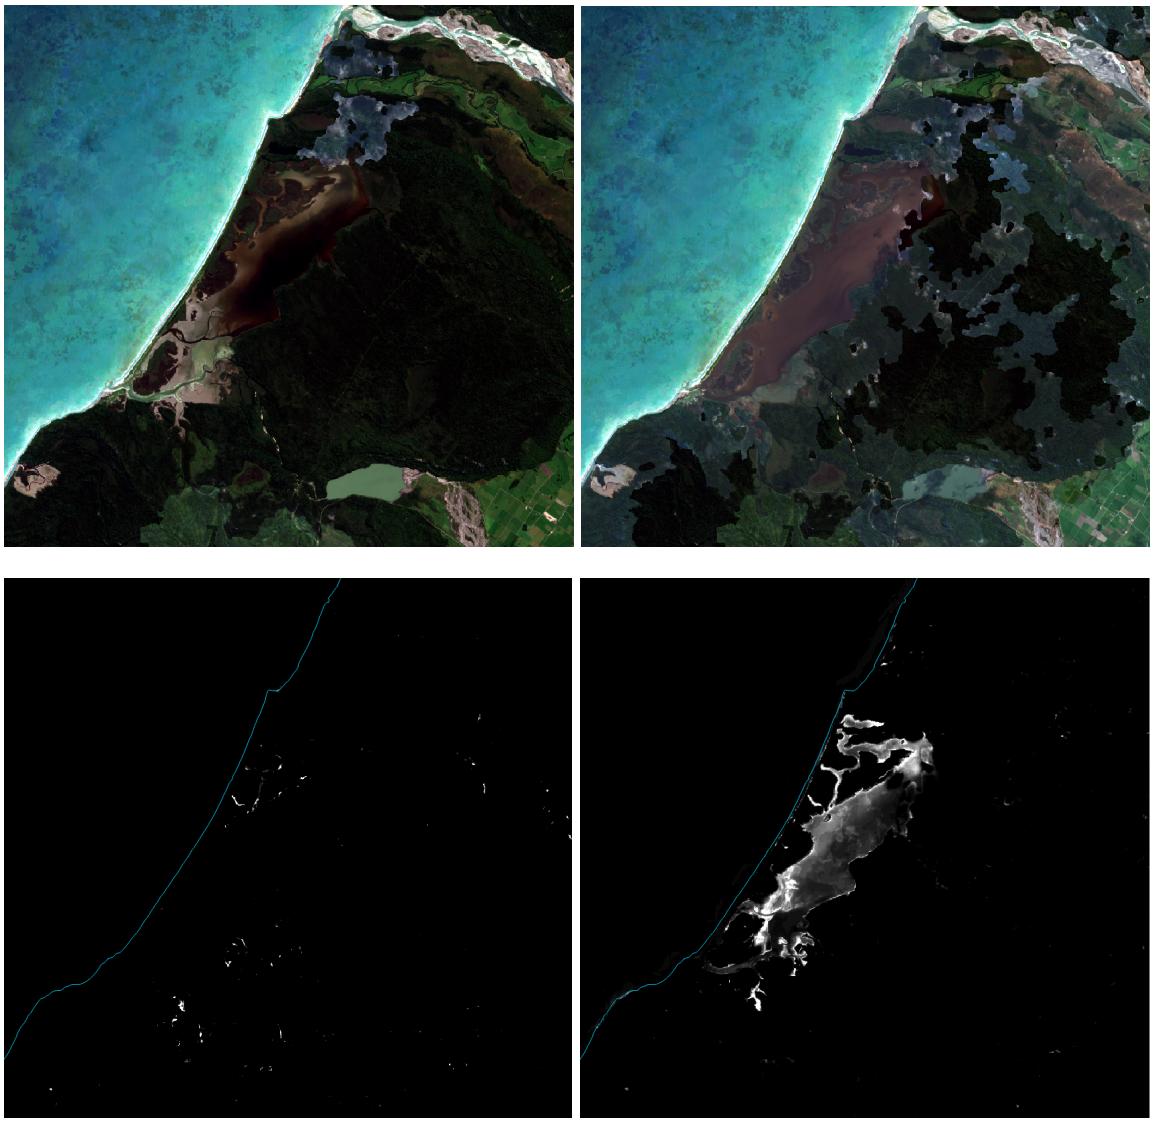

Figure 31. Estuary, West Coast (NZTM 1375151, 5213414). Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

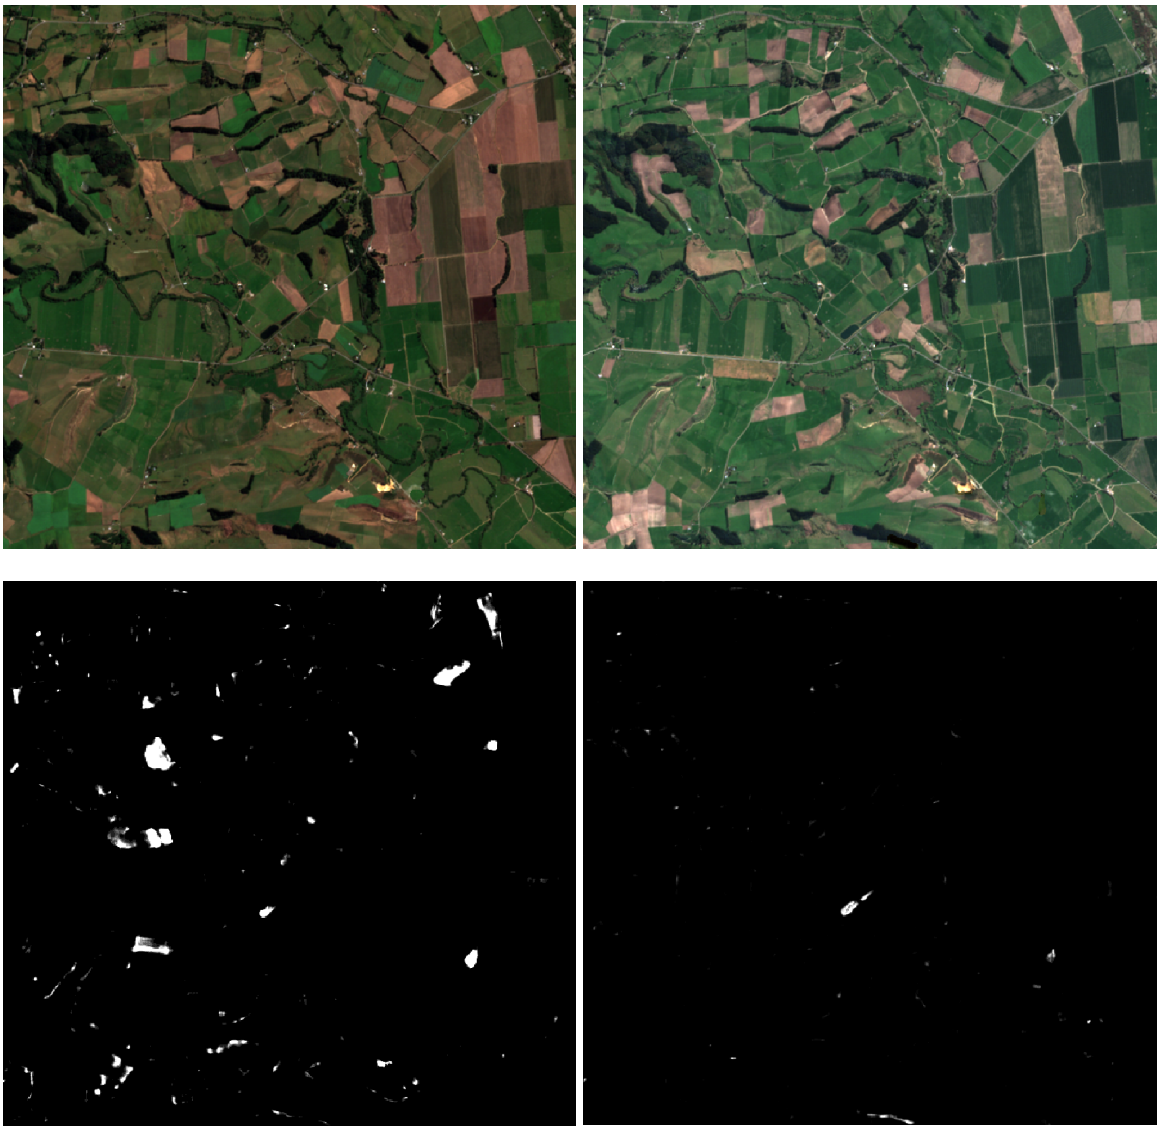

Figure 32. Agricultural land, Mid-Canterbury Plains (NZTM 1453586, 5111179). Top: Sentinel-2 images for March (left) and September (right) 2020. Bottom: predictions for the two-image model (left) and combined model (right).

Both models produce false positives for the non-alpine areas, but this is much less pronounced for the combined model: because it was trained on a combination of visual and change and wetness variability, it tends to pick up spurious water bodies, whereas the two-image model also falsely detects areas of other significant visual change, such as changes in management of forestry and agricultural land. In both cases, at least some of the false positives could be removed by filtering the raster by land cover / land-use type.

5.3 Detecting pastoral wetlands using deep learning

Detecting pastoral wetlands in the Whangape catchment

We trained and tested two models using imagery and wetland polygon data for the Whangape catchment in Waikato: a 'two-image' model consisting of 10-band images for February and October, and a 'combined' model that included 12 monthly wetness bands. The imagery was tiled into 108 256 × 256 pixel images, of which 54 were used for training, 10 for validation during training, and the remaining 44 held out for testing.

There was very little difference in performance between the two models. In hindsight this could have been expected, because there was very little information in the wetness bands for this particular catchment (i.e. wetness was mostly zero, apart from for larger open-water bodies). The models both converged very quickly, with the combined model reaching minimum validation loss after only 16 epochs because of the small training set size. Figure 33 shows the accuracy and loss curves for the training and validation sets during training.

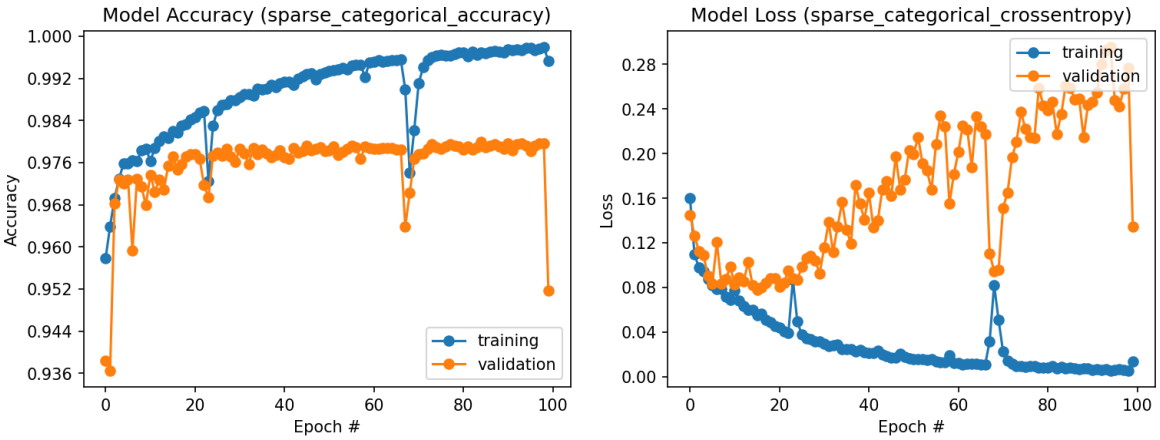

Figure 33. Model accuracy (left) and loss (right) as training progresses. Blue = training set accuracy/loss, orange = validation set accuracy/loss.

The model classified the wetlands in the test set with a per-pixel recall of 72% and a precision of 91%. As for the kettle-hole wetlands, the class prediction substantially under-reports small wetlands because of the high bias in the data (i.e. only around 6% of pixels are classified as 'wetland'). However, inspecting the probability raster shows that the model does pick up smaller wetlands, albeit with fairly low probability. Figure 34 shows the probability raster generated by the combined model for the Whangape catchment.

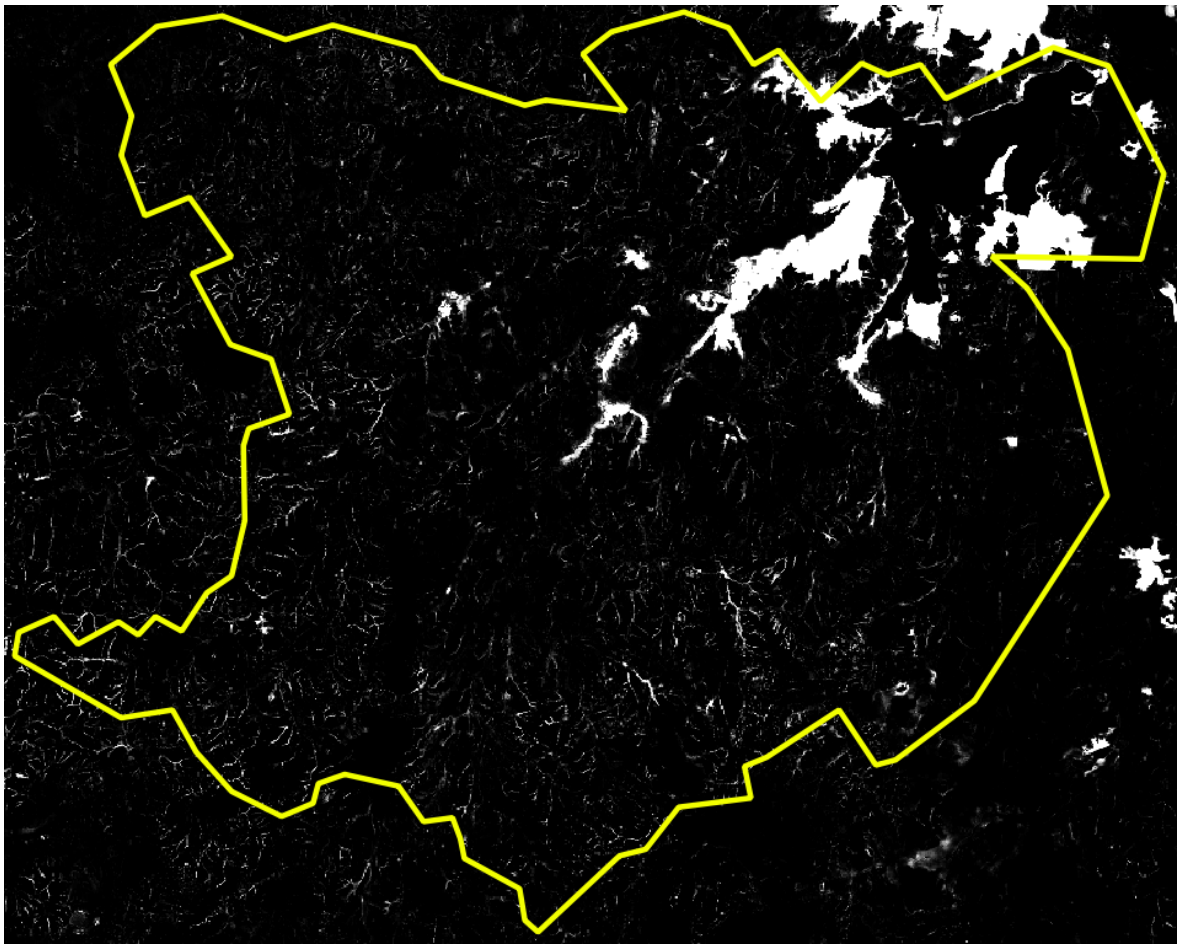

Figure 34. Pastoral wetland probability raster for the combined model after 16 epochs. Yellow line shows the approximate boundary of the Whangape catchment. The intensity of the white areas denotes the probability that this area is a wetland (brighter = higher probability).

We also tested whether stopping at the epoch with the lowest loss meant training was being stopped too early. We generated a probability raster using the model after 69 epochs of training: at this point the model's loss fell to almost as low as the lowest loss (at epoch 16). However, this model returned a substantially worse probability raster, with a high level of false positives. We therefore concluded that using the earlier model with the lowest validation loss was appropriate.

Figure 35 compares the probability raster for the 16-epoch model to 2012 0.5 m aerial imagery for a sample of the Whangape catchment test images.

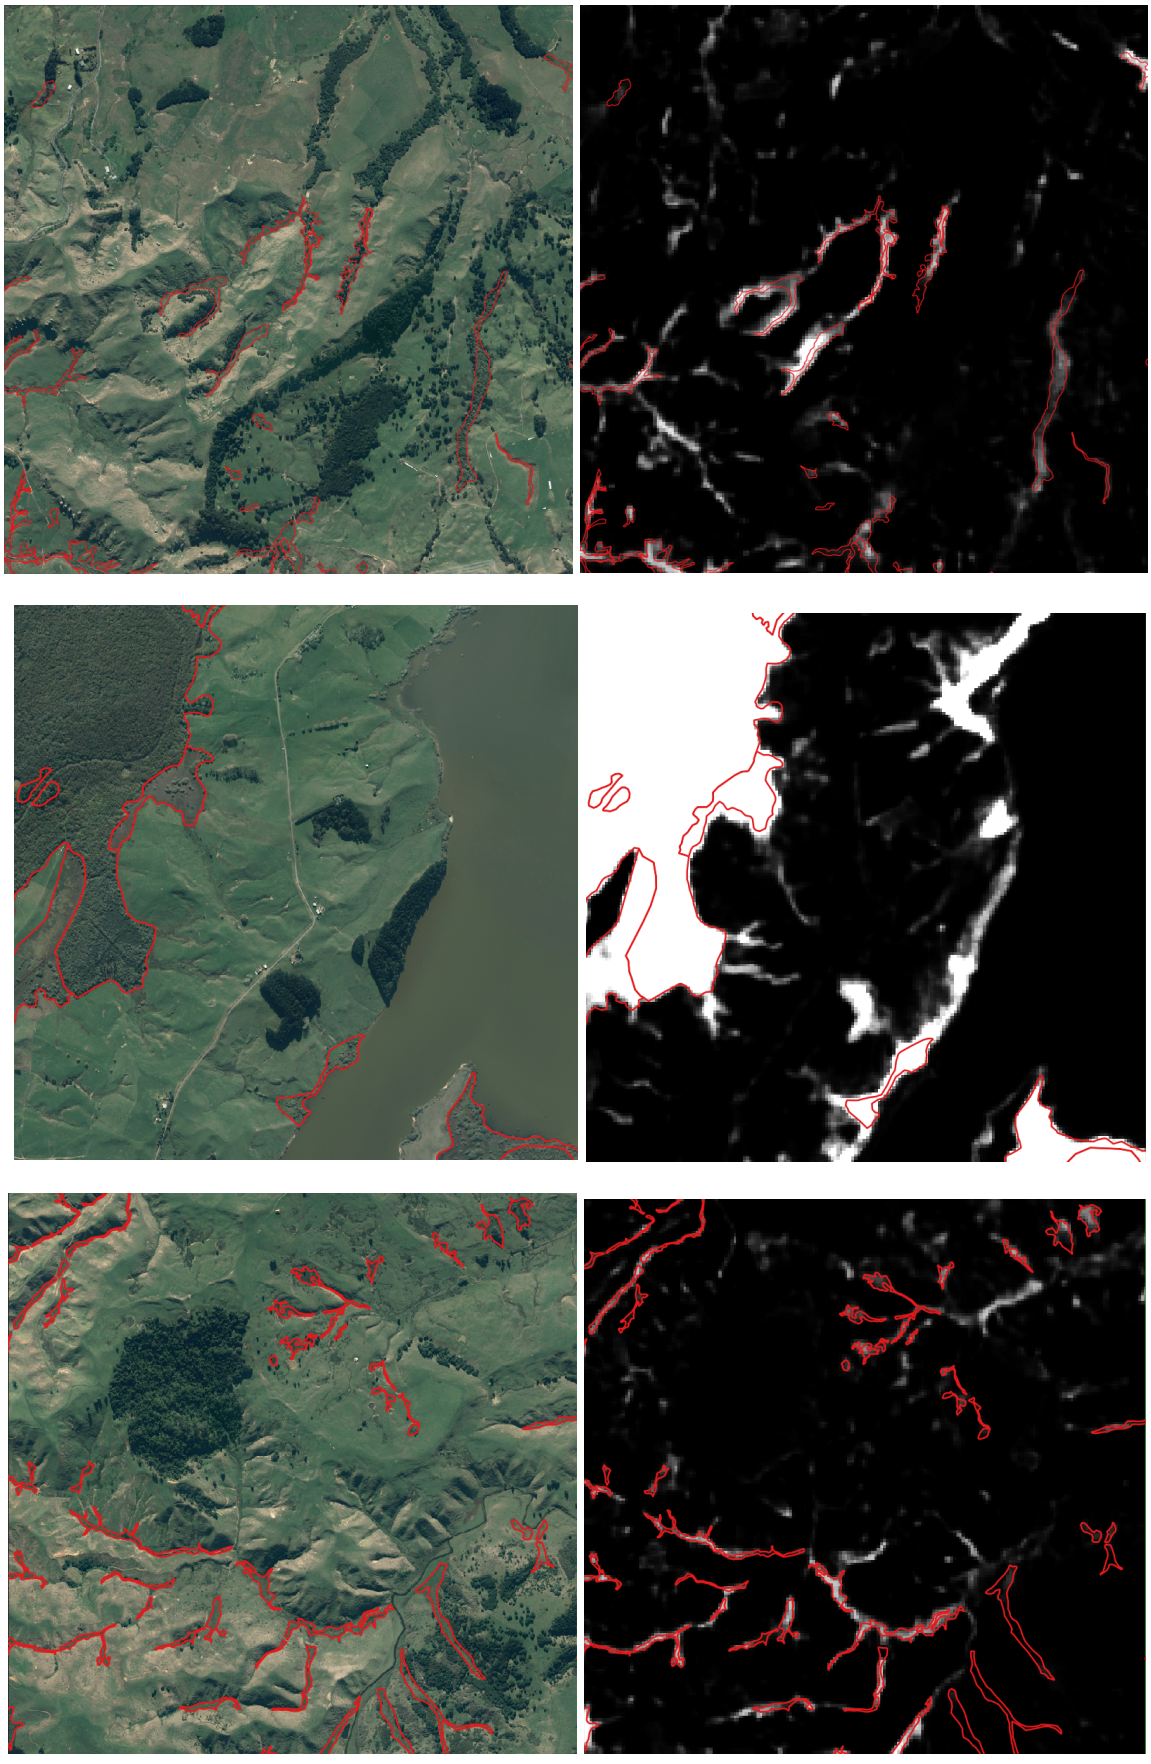

Figure 35. Examples of test areas: aerial imagery (left) and probability raster from the 16-epoch combined model (right). Red polygons are mapped wetlands.

Overall, the model appears to detect wetlands in the Whangape catchment quite well. Figure 36 shows some examples where the model disagrees with the Whangape mapping, particularly false positives.

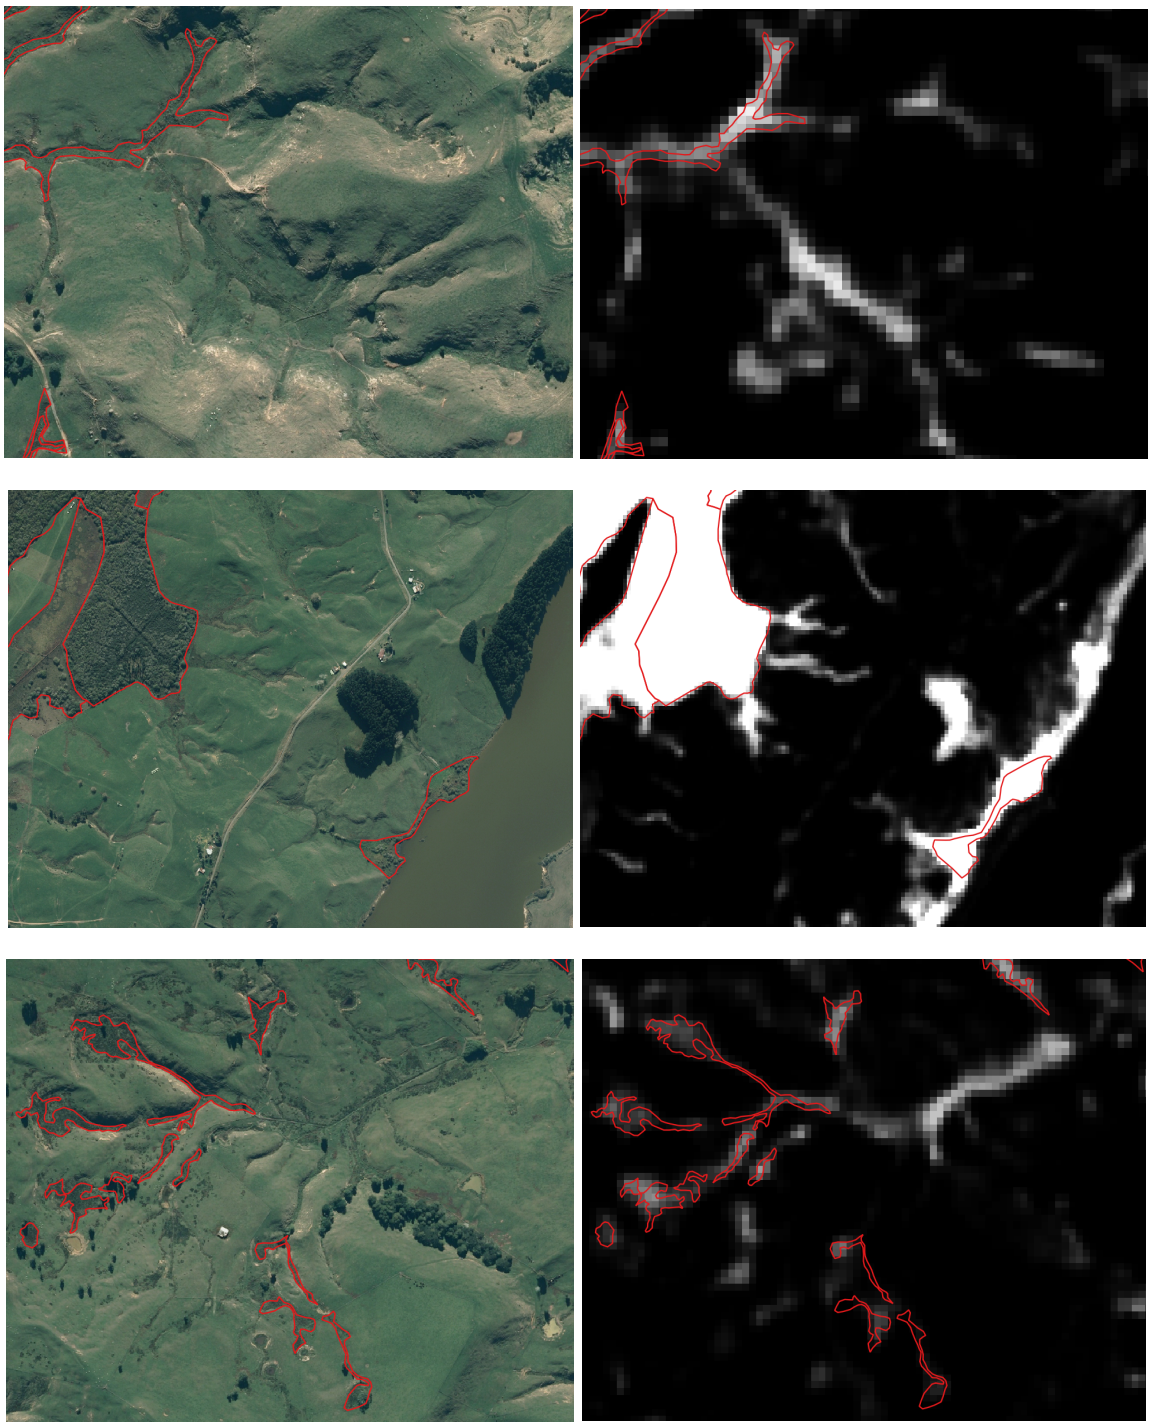

Figure 36. Closer view of apparent false positives and false negatives for the parts of the areas in the previous figure. Red polygons are mapped wetlands.

Notes: In all three cases, apparent false positives (white areas outside the mapped polygons) appear in depressions that have the potential to be wetlands. The middle raster also contains false positives in a forest, as well as at the shore of a water body where the water line varies with time. Some very narrow wetlands in the bottom example appear to have not been detected (false negatives).

We make the following observations.

- Many of the apparent false positives are vegetated depressions that could potentially be wetlands; others are the changing shoreline of open-water bodies.
- Forested/wooded areas are sometimes picked up as false positives.
- The model generally over-reports wetlands; this could be controlled by filtering the raster to include only pixels with values greater than a certain threshold.

Detecting pastoral wetlands across Waikato

We used the 16-epoch combined model to generate a prediction across the entire Waikato region. Because the model was trained for pastoral wetlands only, we used LCDB5 to filter the map, leaving only areas of the following classes that may contain pastoral wetlands: high-producing exotic grasses (class 40), low-producing grassland (class 41), tall tussock grassland (class 43), depleted grassland (class 44), herbaceous freshwater vegetation (class 45), and deciduous hardwoods (class 68). Figure 37 compares the prediction raster before and after filtering.

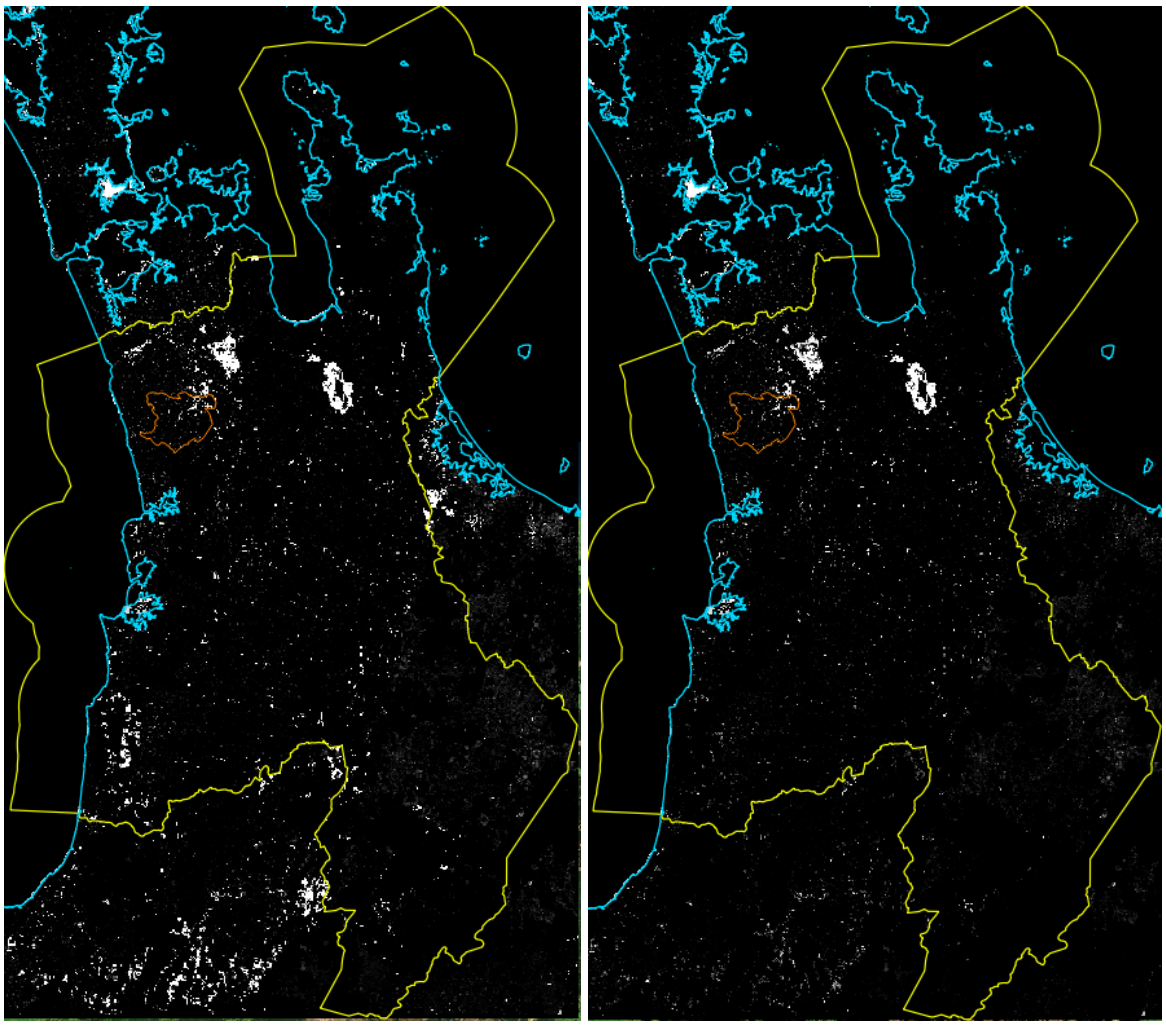

Figure 37. Pastoral wetland probability raster from the 16-epoch combined model for the Waikato region (yellow polygon): raw raster (left) and filtered (right). The small orange polygon denotes the approximate Whangape catchment boundary used for training.

Figure 38 compares the probability raster with Sentinel-2 imagery for a sample of other pastoral areas outside the Whangape catchment where wetlands have also been extensively mapped.

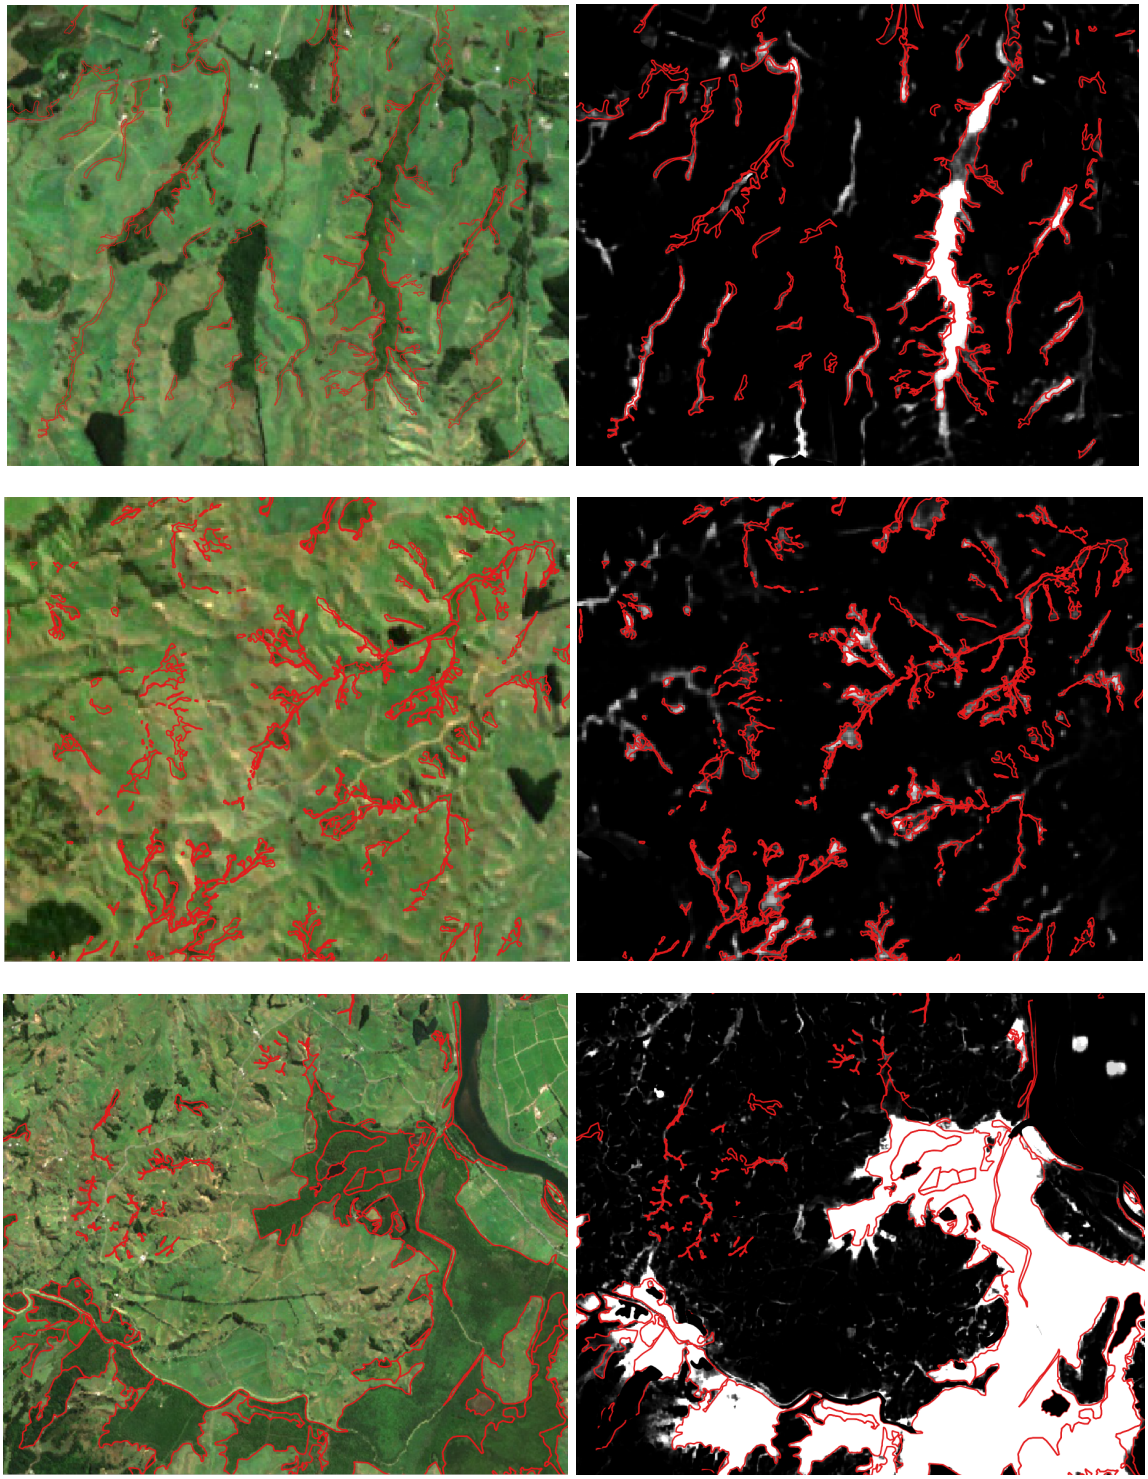

Figure 38. Sentinel-2 imagery and filtered 16-epoch combined probability raster for a sample of well-mapped pastoral areas outside the Whangape catchment. Predominant wetland types are swamps and seepages.

The model appears to detect swamps and seepages well, including small wetlands (less than 1 ha in size) and very narrow wetlands (1–2 pixels wide).

Figure 39 compares the probability raster to 2012 0.5 m aerial imagery for a sample of areas in north Waikato that contain apparent false positives, but where the terrain is similar to the Whangape catchment.

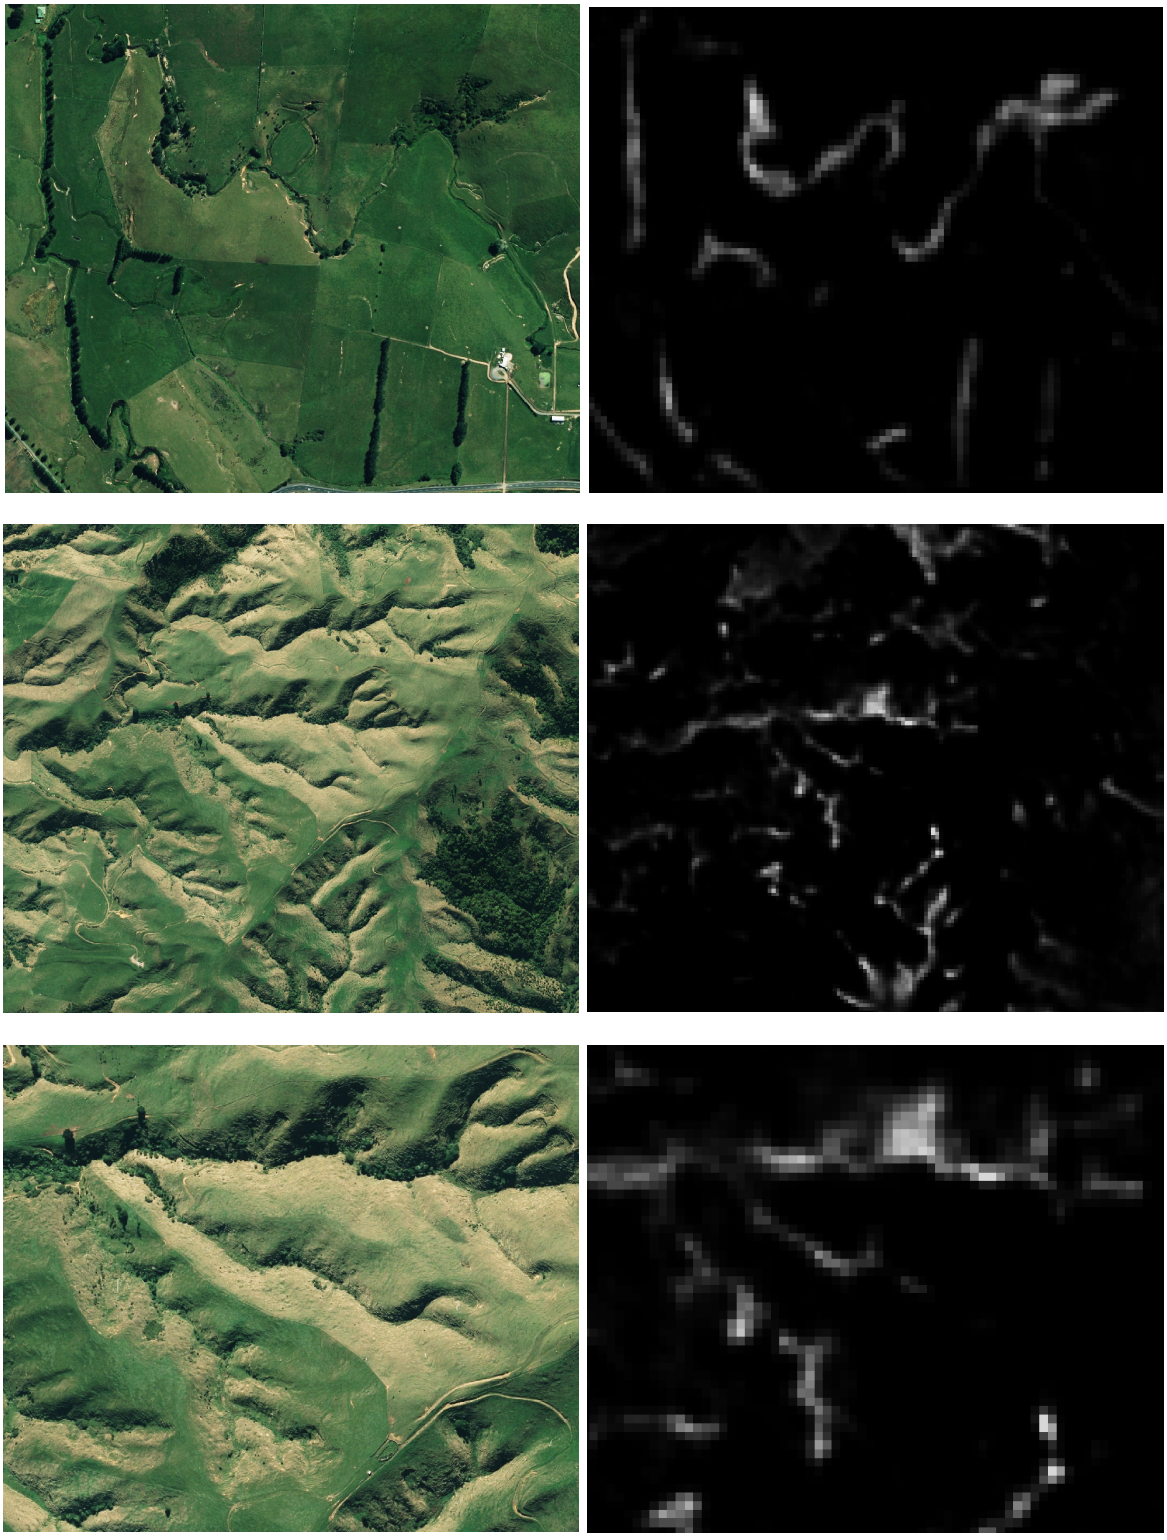

Figure 39. Aerial imagery and probability rasters for sample areas in north Waikato similar to the Whangape catchment.

For the areas with similar topography to Whangape catchment, the model appears to detect plausible wetlands (i.e. vegetated depressions that look similar to the wetlands in the Whangape catchment), although it may also falsely detect wetlands in other vegetated areas. In contrast, Figure 40 shows areas with strong detection in south Waikato, where the terrain is either undulating or more hilly, the (pumice) soil has different drainage properties, or the wetlands are significantly different in appearance (e.g. the swamps and fens south of Lake Taupō).

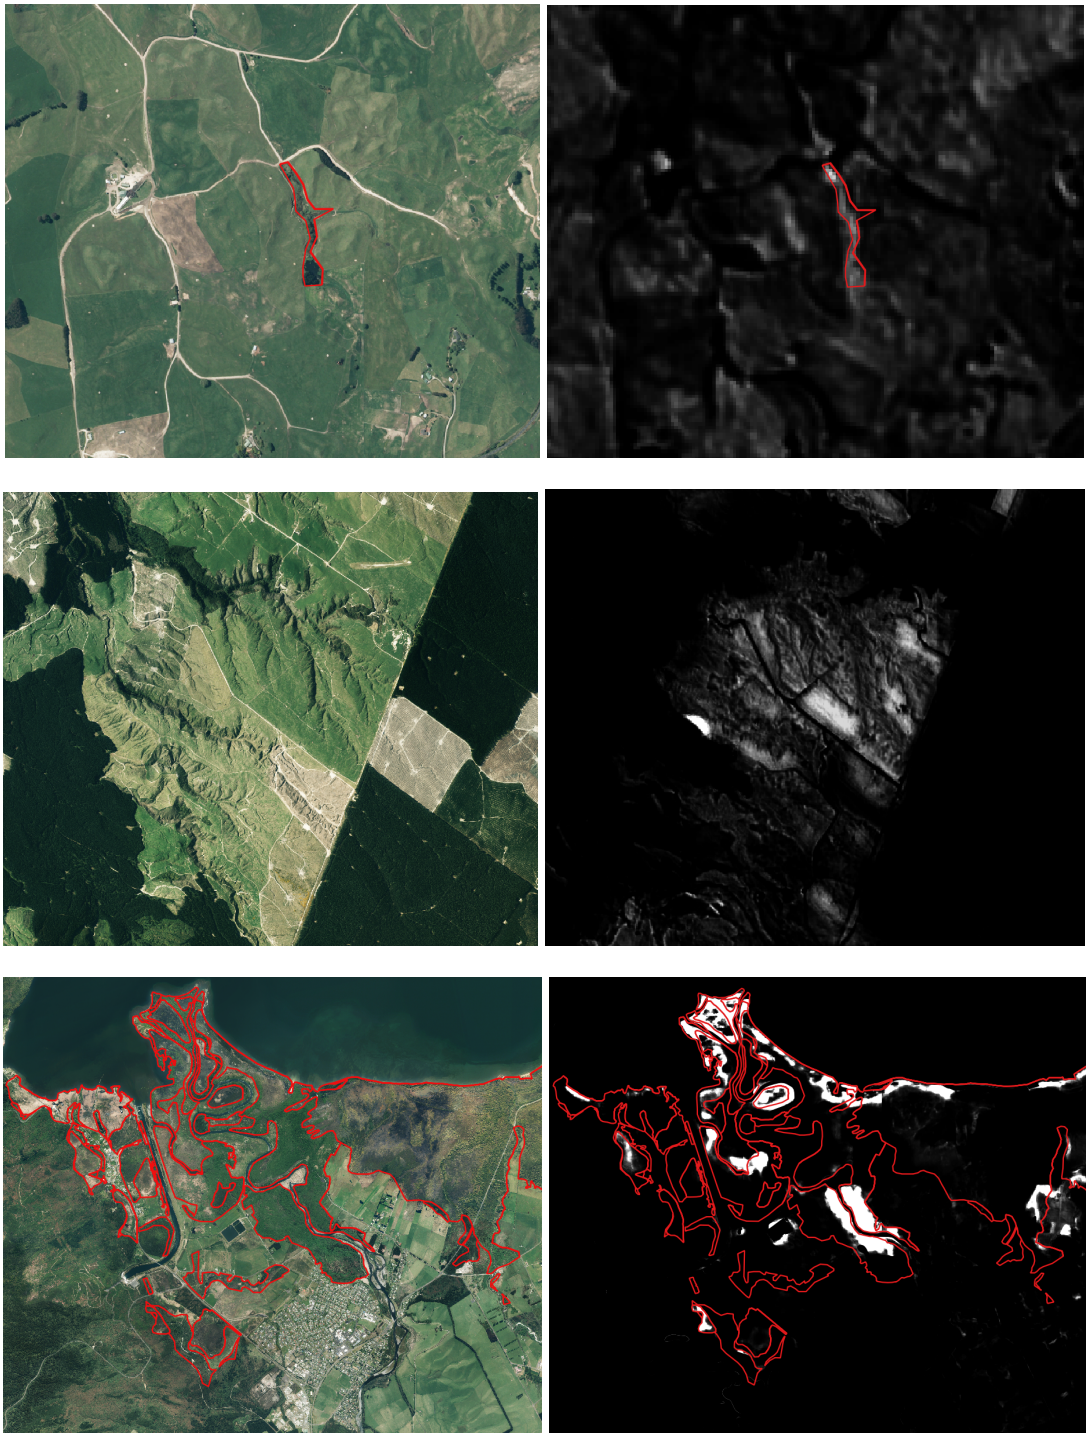

Figure 40. Examples from south Waikato: aerial images (left) and 16-epoch combined probability raster (right).

Prediction for the south Waikato areas is fairly poor, with an increase in both false positives and missed wetlands. This suggests further training is necessary using mapped areas in south Waikato so that the model is exposed to a wider range of terrain and vegetation cover.

Detecting pastoral wetlands outside the Waikato region

We inspected the probability raster for a sample of the Auckland area to assess the potential for the model to extrapolate to areas well outside the training area. As for Waikato, we filtered the result to remove non-pastoral areas. Unlike the Waikato region, Auckland has relatively few pastoral wetlands mapped. Figure 41 shows how the model performs for some examples of pastoral areas in the southern Auckland region.

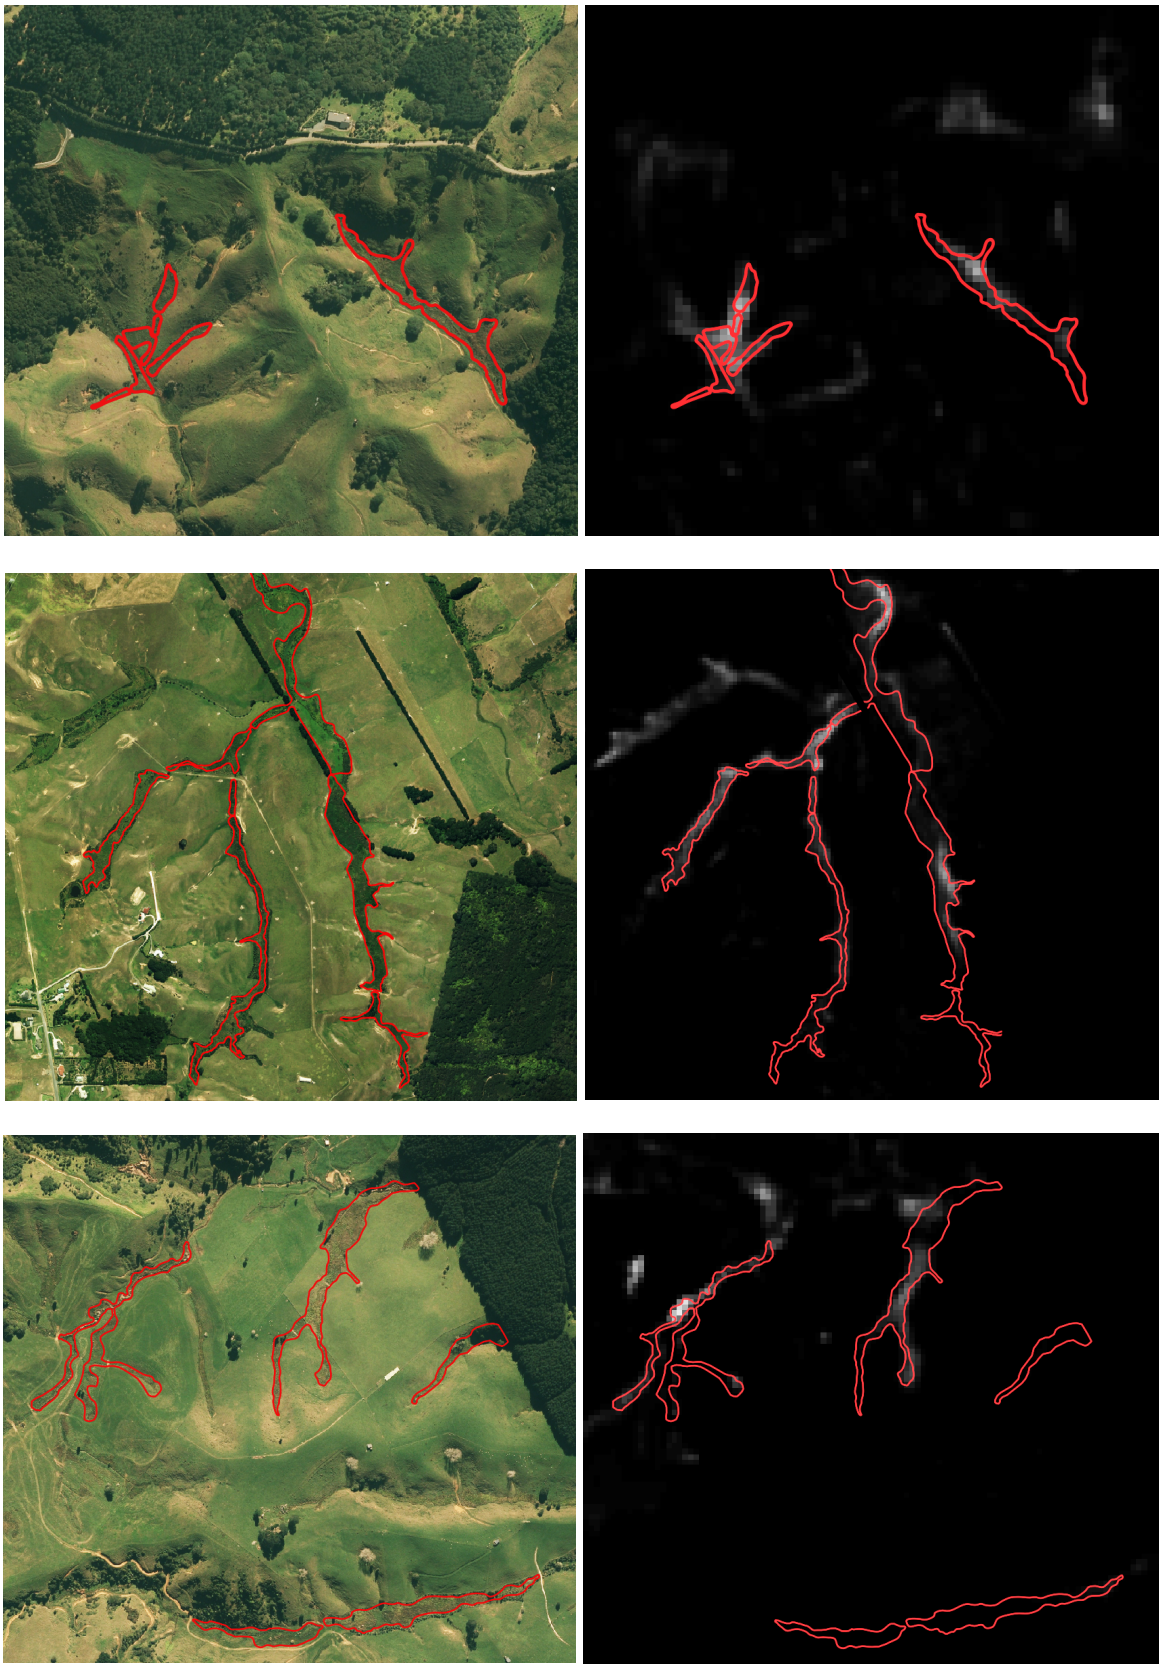

Figure 41. Examples of detection of known wetlands in the Auckland region.

The model detects some – but not all –of the currently mapped wetlands in pasture, and the signal (probability) is generally low. This may be partly because of differences in either the vegetation present in the wetlands or the timing of wet and dry periods. We note that

the visibly driest and wettest months in 2020 were different for Waikato (February and October) and Canterbury (March and September), for example.

The model also detects many (potential) wetlands that are not currently mapped. Figure 42 gives some examples of these.

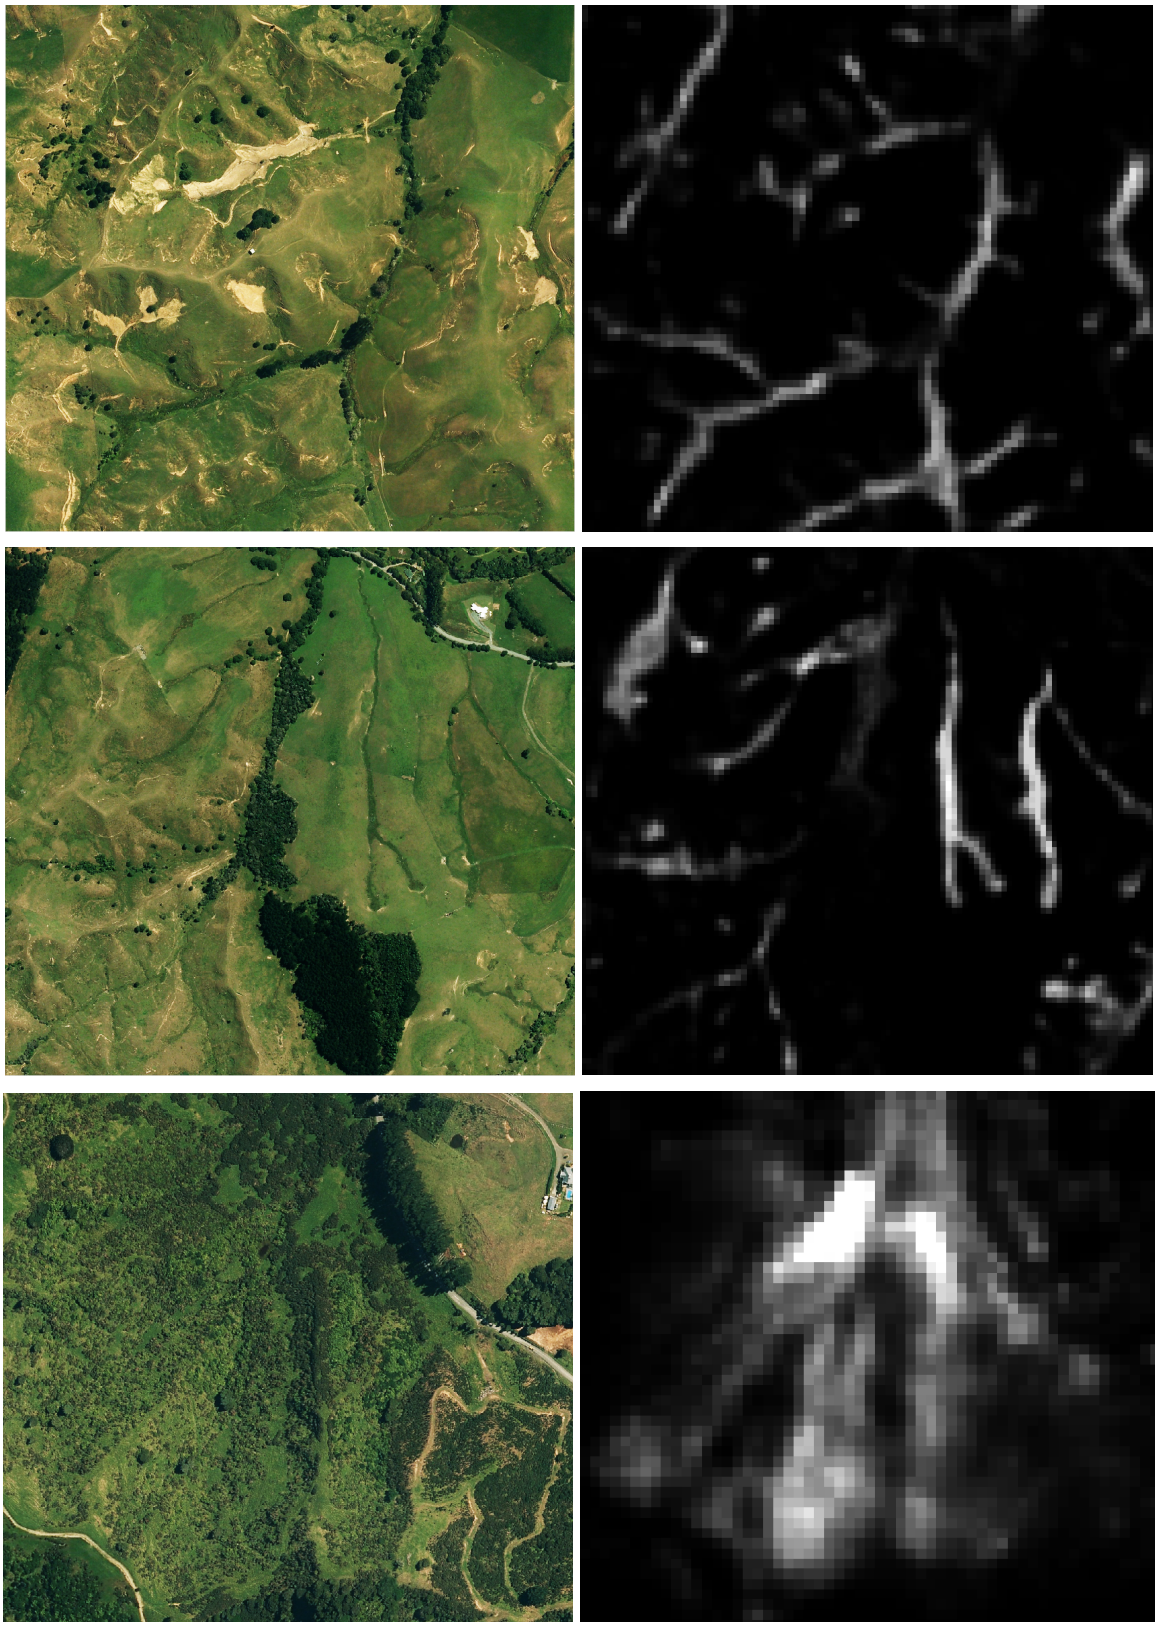

Figure 42. Examples of potential wetland detection in Auckland where none are currently mapped.

Many of the potential wetlands detected appear very similar to those in Waikato (i.e. vegetated valley depressions; Figure 42, top and middle). Others (e.g. Figure 42, bottom) are in quite different terrain. In both cases the results should be verified via ground inspection or a detailed desktop analysis.

6 Guidelines for mapping difficult wetlands

6.1 Mapping ephemeral wetlands using the wetness-based layers

Both the medium-resolution water frequency layer and the high-resolution water mask can be used to locate areas of short-term and permanent water cover. The absolute counts of water observations in conjunction with other environmental information (e.g. timing within a growing season) can be indicative of ephemeral wetlands. Several rule-based systems using expert knowledge, timings, and environmental data (such as excluding permanent water bodies) are conceivable but were not possible to apply in the time available.

The high-resolution mask adds fine-scale water features and can be easily applied to new RGBi imagery when available. The pre-trained model can be either reused or a new model trained with additional training data from new imagery. We recommend that image capture, such as aerial or drone flights, be targeted to times of the year when ephemeral wetlands are most likely to be visible. Unfortunately this is often not when the regional aerial photography in New Zealand is acquired. Buffering around small, semi-permanent water bodies and intersecting them with vegetation information can be useful for further discriminating water body and vegetation types (Figure 43).

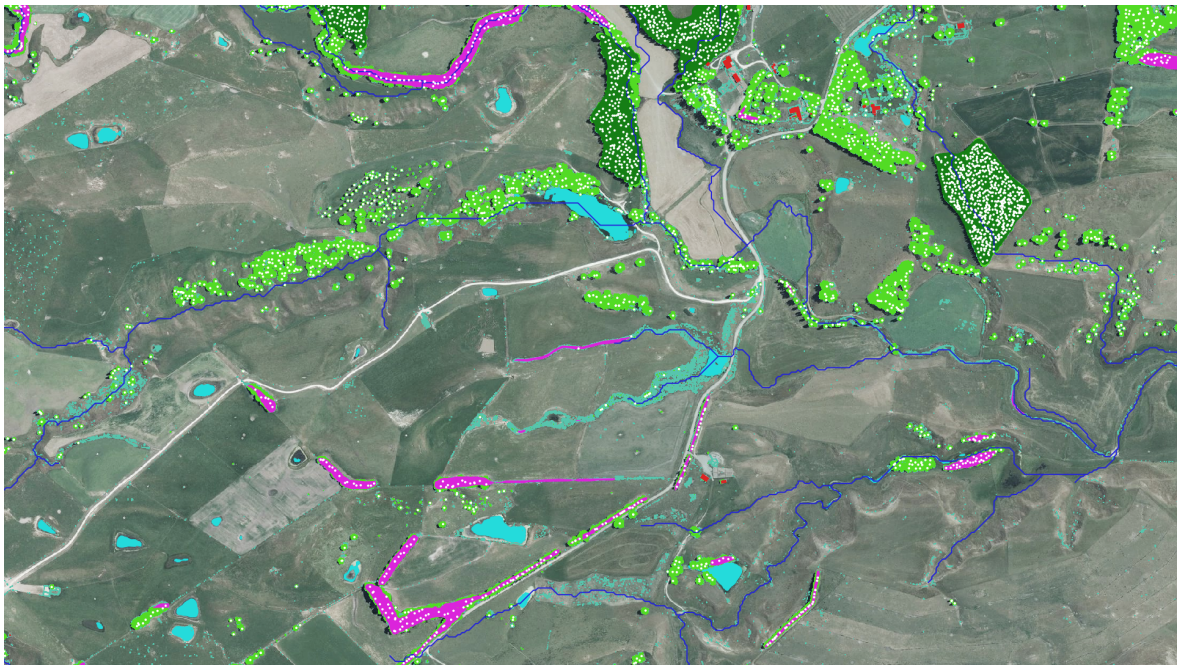

Figure 43. Example of vegetation classification (water bodies in cyan, individual trees as white dots, tree cover in green, pine in dark green, shelter belts in pink and river network in blue). Note: Buffered water area intersected with vegetation types may help categorise water and wetland type.

6.2 Mapping ephemeral wetlands using the deep learning layers

The probability rasters generated by the models can be used directly as a visual aid to determine areas that might contain wetlands of the type that each particular model detects (small water bodies, ephemeral wetlands or pastoral wetlands). For example, the raster can be clipped to an appropriate minimum value and rendered using false colouring to provide a quick visual indication of likely wetlands.

Alternatively, the raster can be converted to polygons by thresholding the probability raster to generate a wetland mask, converting the areas classed as wetlands to polygons, and filtering the resulting layer to remove very small polygons, which are more likely to be noise. Figure 44 shows the polygons generated for the CSIGG where the raster was thresholded (at probability ≥ 10/255), and the resulting polygons were then filtered (to area ≥ 2,000m^2), whereas Figure 45 shows the same polygons for the area east of Lake Pukaki. At this threshold the CSIGG wetland polygon layer contains around 4,000 polygons, which is a sufficiently small number to manually inspect and classify as wetland/not wetland. The accepted polygons could then be further verified and refined to give a final map.

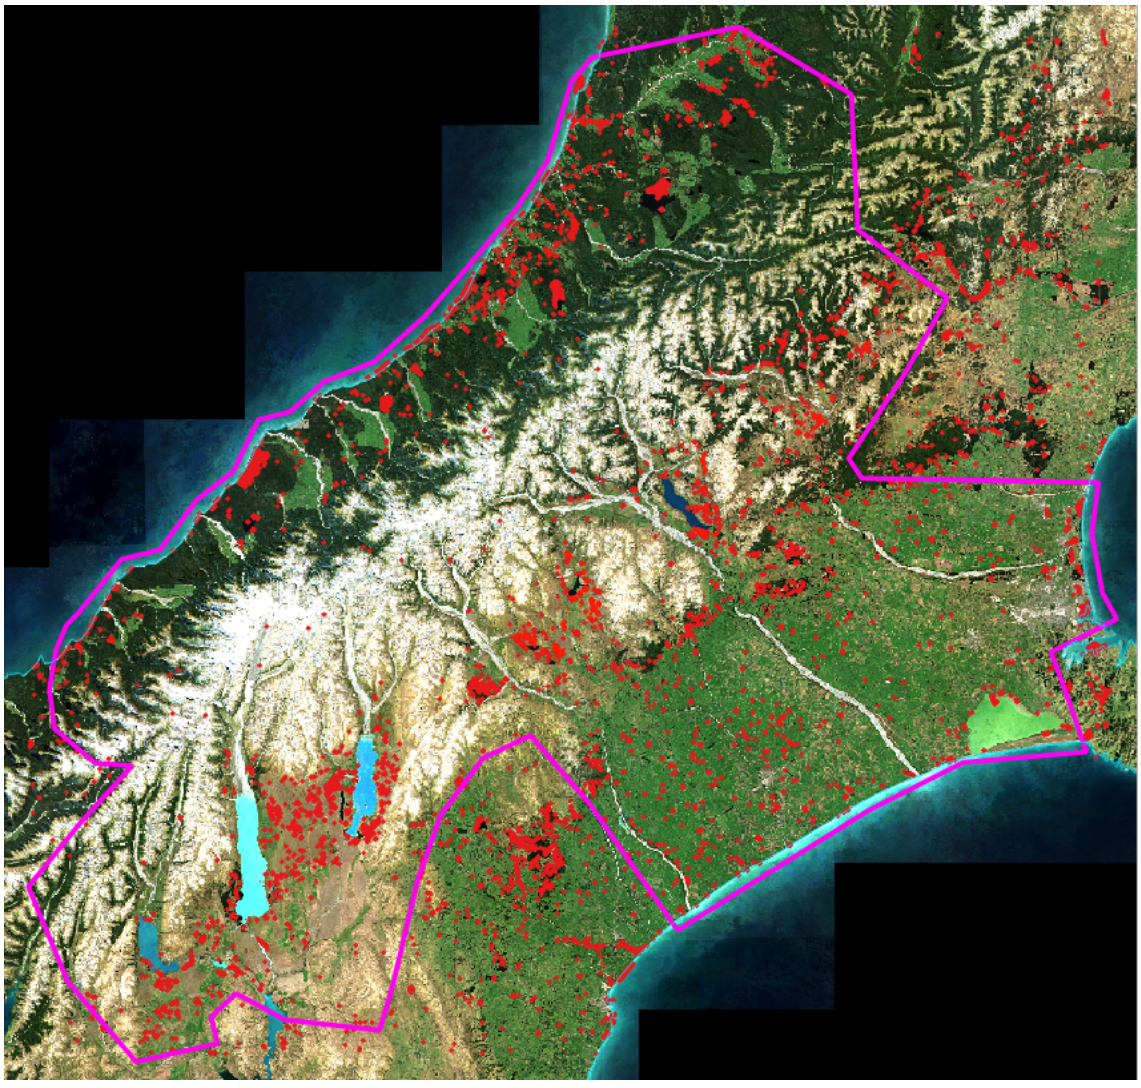

Figure 44. Potential ephemeral wetland polygons for the CSIGG area of size 2,000m^2 or greater.

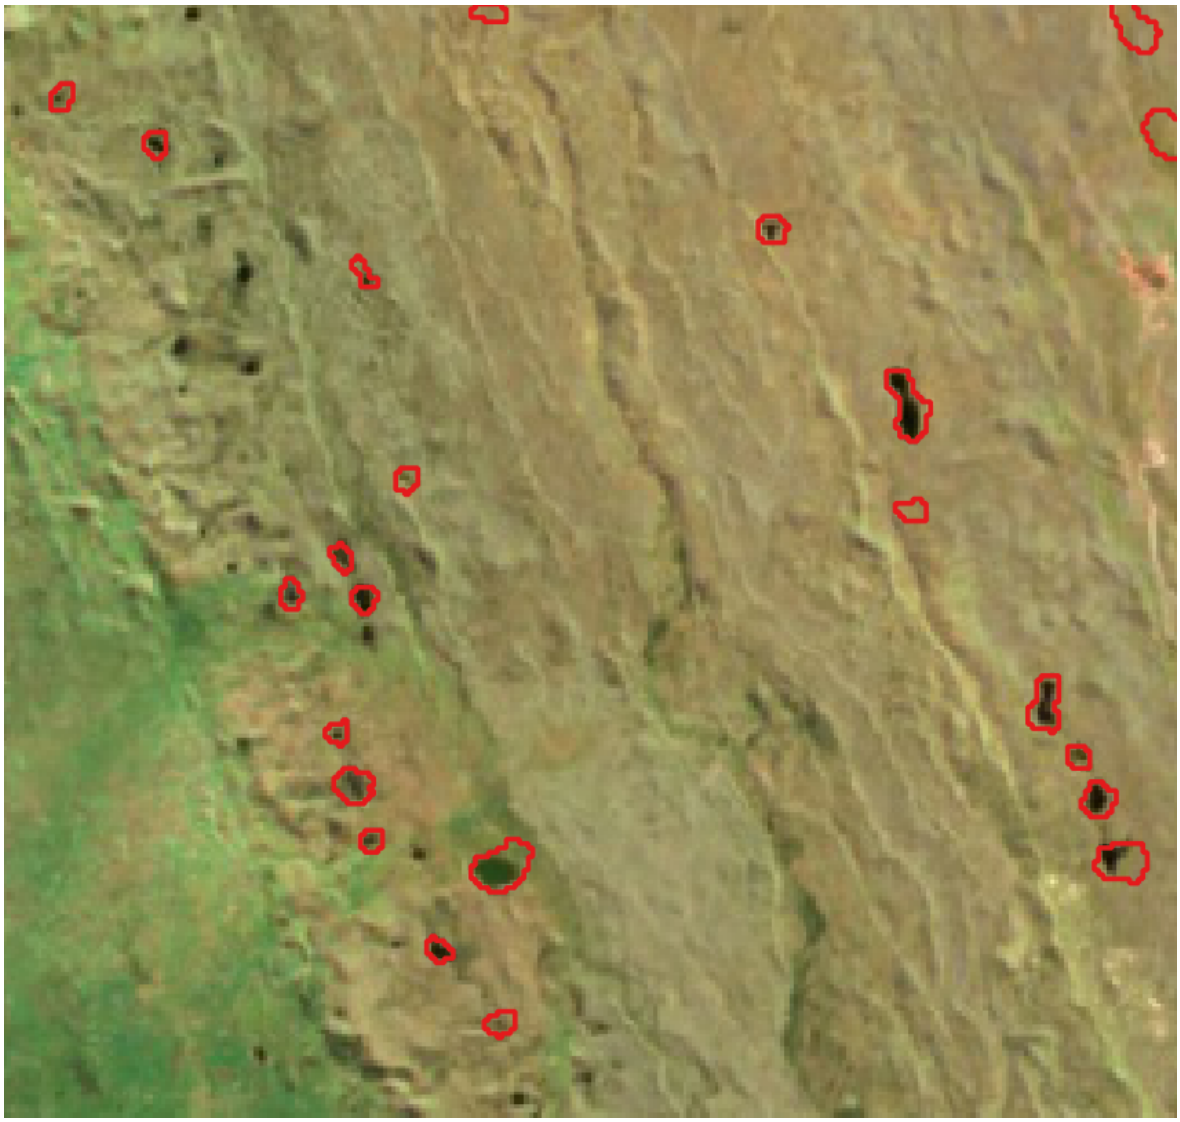

Figure 45. Example of potential kettle-hole polygons east of Lake Pukaki.

6.3 Mapping pastoral wetlands using the deep learning layers

A similar approach can be used to map pastoral wetlands. In Figure 46 the wetland probability raster has again been thresholded (at $p \geq 10/255$) and the resulting polygons filtered (to area ≥ 2,000m^2).

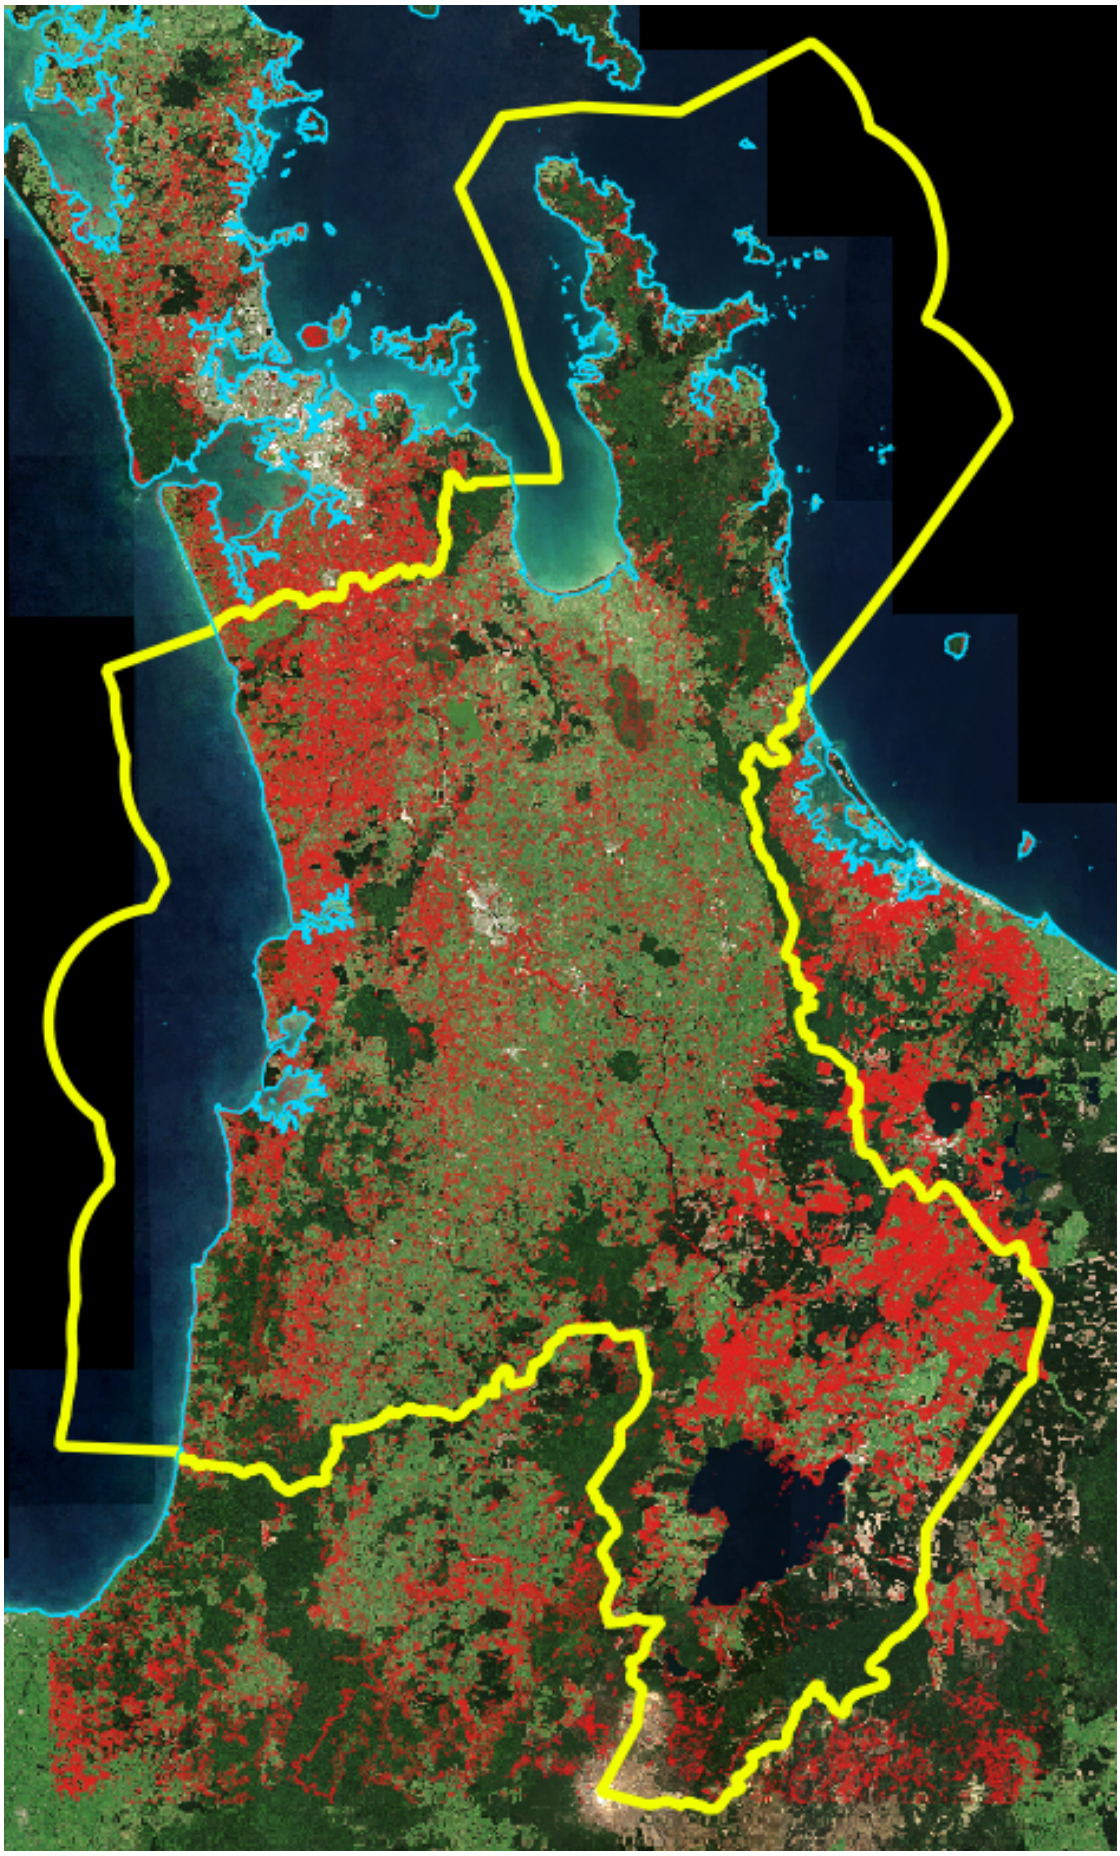

Figure 46. Potential pastoral wetlands in the Waikato region of size 2,000m^2 or greater.

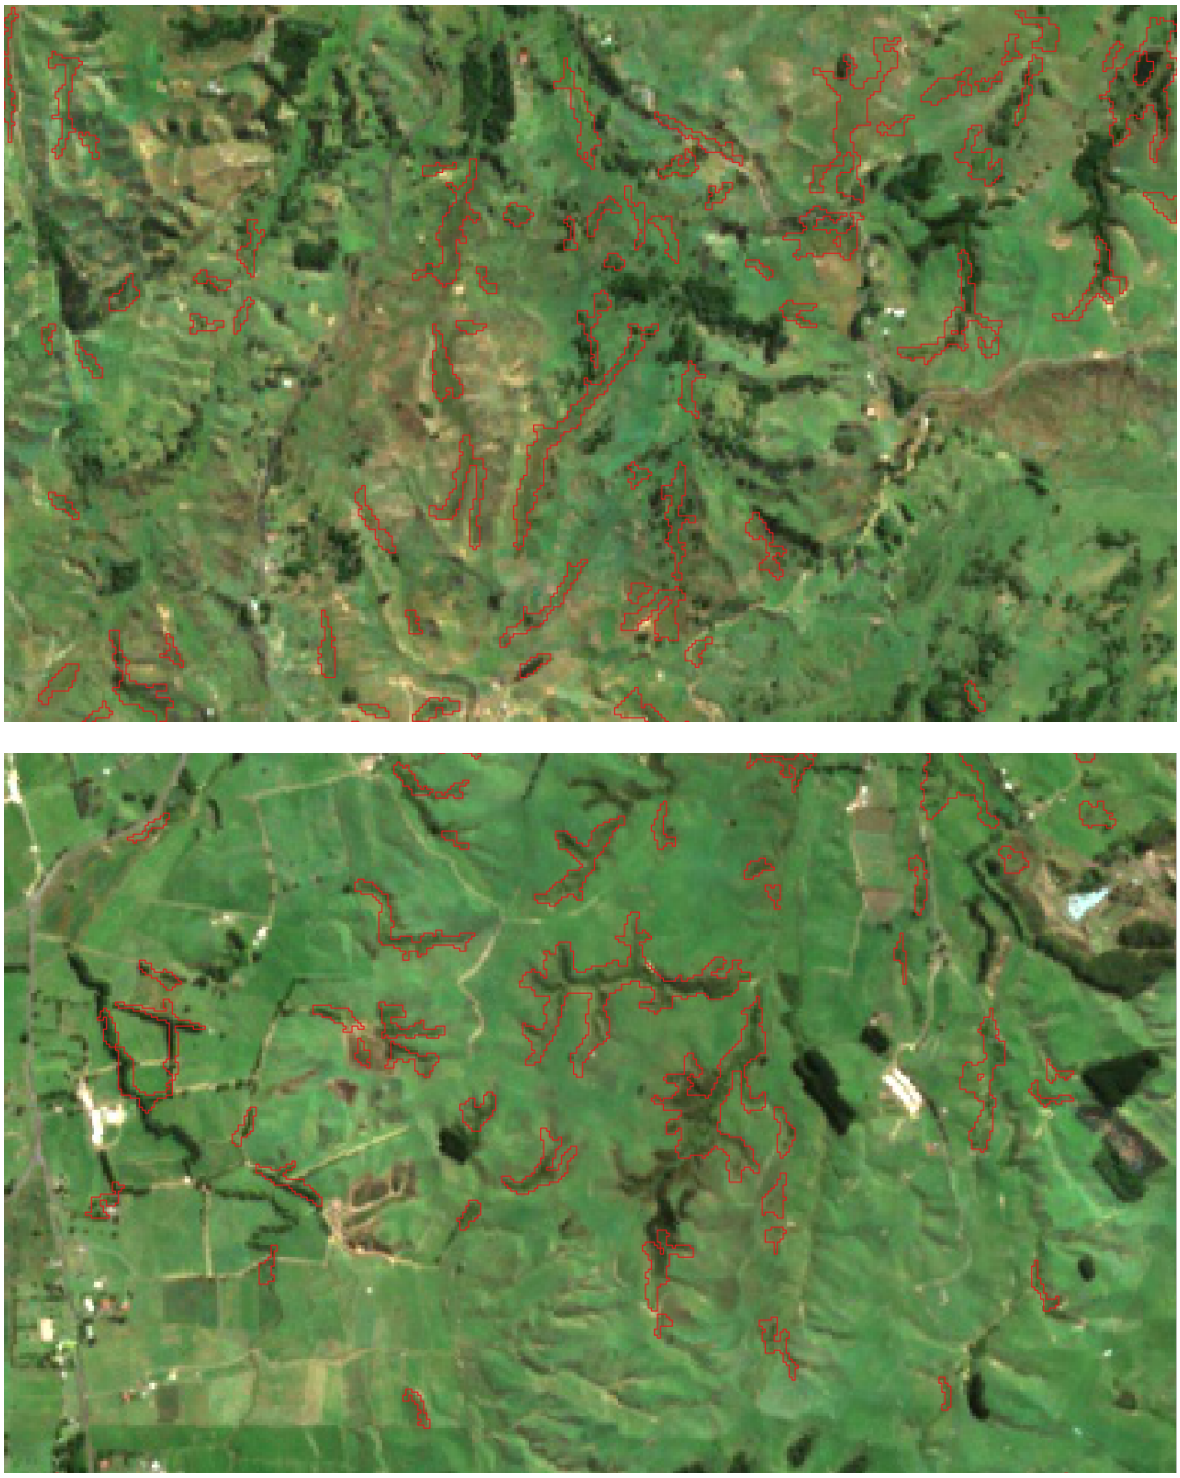

Figure 47. Examples of candidate wetland polygons in hilly areas. Top: north of Whangape catchment (NZTM 1768948, 5860734). Bottom: East of Hamilton (NZTM 1817203, 5821613).

The accuracy of the proposed polygons varies with terrain. For flat, highly cultivated farmland the rate of false positives clearly increases significantly, and it may be prudent to exclude these areas from consideration. In hilly areas, such as in Figure 47, accuracy is likely to be significantly higher. We suggest that the polygons in this area could be used to highlight areas to be investigated and potentially mapped by considering the instructions below.

1 Filter the polygons by size to give a manageable number to check; e.g.:
   - 639,569 unfiltered polygons
   - 285,084 polygons ≥ 500m^2 (NPS-FM requirement)
   - 105,863 polygons ≥ 2,000m^2
   - 47,852 polygons ≥ 5,000m^2
   - 24,600 polygons ≥ 10,000m^2 (1 ha)
2 Quickly scan the polygons and triage them into 'keep', 'discard', and 'review'.
3 Further assess the 'review' polygons and reclassify as many as possible into 'keep' or 'discard'.
4 For the 'keep' polygons, manually map or refine their extent.

7 Conclusions

7.1 Wetness-based detection of potential ephemeral wetlands

Ephemeral wetlands are of high ecological and cultural importance, and it is essential to improve mapping of their detection and extent at national scale in New Zealand. Our study aimed at providing practical methods and data sets that can be utilised by expert practitioners and regional councils.

We developed a transparent remote-sensing method identifying seasonally inundated areas and compared them with ground-truth locations of ephemeral wetlands in Canterbury. Many of these wetlands were successfully identified, and we provided recommendations for how to target cases that are inherently hard to map (i.e. very small, vegetated ponds) from 10 m satellite imagery. We generated a prototype layer that councils can use to create candidate wetland polygons, thereby assisting manual inspection and delineation tasks.

The deep learning-based water detection mask using high-resolution aerial photography provides a valuable resource in addition to the frequency layer. This high-resolution mask adds fine-scale water features to the data set and can be easily applied to new RGBi imagery, when available. It can be used as a mask to constrain the frequency layer and reduce false predictions over forests and other land-cover types.

Both layers – the medium-resolution water frequency and high-resolution water mask – are useful layers beyond wetland applications and proved stable for many different water body types, including braided rivers, streams, ponds, kettle holes, lakes and coast.

7.2 Image-based detection of potential wetlands

This small study investigated whether the mapped kettle-hole wetland areas from McMillan and Wiser were sufficient to train a model to detect ephemeral wetlands in general. The training data set was very small, but the model shows promise, giving reasonable sensitivity with a fairly low false positive rate when applied to areas of similar

geography across the CSIGG. From the probability raster generated by the model we were able to generate a useful polygon layer that could be used to guide mappers to potential ephemeral kettle-hole wetlands in the CSIGG.

We applied the same approach to pastoral wetlands in the Waikato by training a model on the well-mapped Whangape catchment. The resulting probability raster appeared to be reasonably predictive of potential pastoral wetlands and might be a useful additional tool for finding this wetland type. The model also appeared to transfer reasonably well to similar terrain in the Auckland region, demonstrating that it might be useful for other parts of the country that are similar in nature.

The resulting outputs, while not sufficiently accurate for automating the mapping process, are nonetheless potentially useful for directing council staff to where they might look for wetlands. The outputs include some false positives that could be simply masked out since they are in areas where the geography does not lend itself to this type of wetland; and others, such as other ephemeral and permanent wetlands, that may be otherwise useful to the mapping process.

8 Recommendations

- The research products described have been visually analysed my MWLR but they have not been independently assessed. We recommend an assessment of the accuracy and usefulness of the products in collaboration with council staff, including an assessment of which (if any) false positives (i.e. detected areas that are not kettle hole ephemeral wetlands (in the CSIGG) or pastoral wetlands (in the Waikato region) are useful for detecting and mapping other types of wetlands.

8.1 Wetness-based methods

It is challenging to capture seasonal, short-term inundated ephemeral wetlands in one-off aerial or satellite images because of timing, weather conditions (cloud cover), and resolution. The Sentinel-2-derived water frequency layer is able to address two of these issues, and even demonstrated the ability to locate wetlands with sizes of less than 0.1 ha. Not all ground-truthed polygons were picked up by the automated algorithm, and we do not know whether this is due to misclassification in some of the images, no actual inundated areas present in the image time series, or misalignment of the ground-truth capture date and Sentinel-2 observations in 2021 (some ephemeral wetlands might have been lost since their observation in 2016).

We also identified several shortcomings of the current water frequency layer. First, small features may not be visible in the 10 m imagery or are spectrally mixed with surrounding vegetation. More focused training that includes these small areas and mixed pixels leads to better detection rates, but at the cost of more false predictions over forested areas. This may be acceptable if the human GIS analyst is trying to map ephemeral wetlands outside forests.

It is often desirable to work with definite classes rather than absolute water frequencies (counts per month or year), which are locally biased due to Sentinel-2 scanning patterns, uneven persistent cloud cover, differing levels of atmospheric aerosols, and other factors. It is possible to turn the water frequency layer into three classes (e.g. permanent water, seasonal, or one-off/flooding events). Calibrating these classes for different water body types seems practical but will need to be carefully done, incorporating local knowledge. Other auxiliary data sets such as environmental factors and indications of the regional growing season are likely to be needed as well.

We therefore recommend the following extensions:

- Extend the water frequency layer to all Sentinel-2 satellite imagery for the period 2016–2024.
- Experiment with turning the water frequency layer into a classification map (permanent water, seasonal wetness, one-off/flooding events).

8.2 Image-based methods

For this approach we used a combination of summer and winter Sentinel-2 10-band imagery and monthly wetness rasters. There is potential to further improve the modelling for both ephemeral and pastoral wetlands. We recommend the following.

- Extend the deep learning training areas to include a greater range of land cover and terrain.
- Add further inputs to the deep learning models, including aerial imagery and synthetic-aperture radar (SAR).
- Increase the deep learning training data time series by adding more seasonal images within the year studied, and potentially multi-year imagery, where possible.
- Carry out further tuning of the model training processes, including trialling alternative methods of modelling temporal change.

9 Acknowledgements

We gratefully acknowledge the Ministry for the Environment for funding this research and the Department of Conservation for instigating previous research in this area, including providing valuable wetland mapping data. We also thank Stella Belliss for reviewing this report, and Ray Prebble for editing.

10 References

Ausseil A-GE, Chadderton WL, Gerbeaux P, Stephens RTT, Leathwick JR 2011. Applying systematic conservation planning principles to palustrine and inland saline wetlands of New Zealand. Freshwater Biology 56(1): 142–161. https://doi.org/10.1111/j.1365-2427.2010.02412.x

Barrell DJA, Andersen BG, Denton GH, Smith Lyttle B 2013 Glacial geomorphology of the central South Island, New Zealand – digital data. GNS Science monograph 27a Geographic Information System digital data files + explanatory notes. Lower Hutt, GNS Science.

Belliss SE, Pairman D, Dymond J, Amies A, Zoerner J, Shepherd J, et al. 2019. Identification of high-risk agricultural activities: national mapping of the location, scale and extent of winter forage cropping and intensive grazing on hill country land. Landcare Research contract report LC3469 prepared for Ministry for the Environment.

Bunting P, Gillingham S 2013. The KEA image file format. Computers & Geosciences 57: 54–58.

Clarkson BR 2014. A vegetation tool for wetland delineation in New Zealand. Landcare Research contract report LC1793.

Cox DR 1958. The regression analysis of binary sequences. Journal of the Royal Statistical Society Series B: Statistical Methodology 20(2) :215–232.

Drusch M, Del Bello U, Carlier S, Colin O, Fernandez V, Gascon F, et al. 2012. Sentinel-2: ESA's optical high-resolution mission for GMES operational services. Remote Sensing of Environment 120: 25–36.

Dymond JR, Shepherd JD 2004. The spatial distribution of indigenous forest and its composition in the Wellington Region, New Zealand, from ETM+ satellite imagery. Remote Sensing of Environment 90: 116–125.

Fraser S, Singleton P, Clarkson B 2018. Hydric soils – field identification guide. Manaaki Whenua – Landcare Research contract report LC3233.

Hosmer Jr DW, Lemeshow S, Sturdivant RX 2013. Applied logistic regression. John Wiley & Sons.

Johnson P 2003. Canterbury kettles and other ephemeral wetlands: centres for native plant diversity. Canterbury Botanical Society Journal 37: 42–47.

Johnson P, Gerbeaux P 2004. Wetland types in New Zealand. Wellington, Astra Print.

McMillan AMS, Wiser SK 2019. Using satellite imagery to detect and map ephemeral wetlands. Manaaki Whenua – Landcare Research contract report LC3555 for Department of Conservation.

MfE 2021. Wetland delineation hydrology tool for Aotearoa New Zealand. Wellington, Ministry for the Environment.

MfE. 2022. Wetland delineation protocols. Wellington, Ministry for the Environment.

MfE 2024. Wetland delineation using desktop methods: a guide. Wellington, Ministry for the Environment.

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. 2011. Scikit-learn: machine learning in Python. Journal of Machine Learning Research 12: 2825–2830.

Qiu S, Zhu Z, He, B. 2019. Fmask 4.0: Improved cloud and cloud shadow detection in Landsats 4–8 and Sentinel-2 imagery. Remote Sensing of Environment 231(15).

Ronneberger OL, Fischer P, Brox T 2015. U-Net: convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-assisted Intervention. Cham, Switzerland, Springer. Pp. 234–241.

Shepherd JD, Schindler J, Dymond JR 2020. Automated mosaicking of Sentinel-2 satellite imagery. Remote Sensing 12: 3680.

Zhu Z, Wang S, Woodcock CE 2015. Improvement and expansion of the fmask algorithm: cloud, cloud shadow, and snow detection for Landsats 4–7, 8, and Sentinel 2 images. Remote Sensing of Environment 159: 269–277.

Zhu Z, Woodcock CE 2012. Object-based cloud and cloud shadow detection in Landsat imagery. Remote Sensing of Environment 118: 83–94.

Zhu Z, Woodcock CE 2014. Continuous change detection and classification of land cover using all available Landsat data. Remote Sensing of Environment 144: 152–171.