Mapping naturally uncommon ecosystems – summary of data and methods

Envirolink Grant: C09X2301

September 2025

Adrian Monks, Gretchen Brownstein, Rowan Buxton, Sarah Richardson, Jan Schindler, James Shepherd, Susan Wiser

Manaaki Whenua – Landcare Research, a group of the Bioeconomy Science Institute

Contract Report registration number: 2526-0032

Prepared for: New Zealand's regional councils

Disclaimer

This report has been prepared by the New Zealand Institute for Bioeconomy Science Ltd for New Zealand's regional councils. If used by other parties, no warranty or representation is given as to its accuracy and no liability is accepted for loss or damage arising directly or indirectly from reliance on the information in it.

© New Zealand's regional councils 2025

This report has been prepared by New Zealand Institute for Bioeconomy Ltd for New Zealand's regional councils and New Zealand Institute for Bioeconomy Ltd has agreed that New Zealand's regional councils owns the copyright in the report. It may not be reproduced or copied, in whole or in part, in any form or by any means without the written permission of New Zealand's regional councils.

Reviewed by:

Jo Carpenter
Researcher
Manaaki Whenua – Landcare Research

Approved for release by:

Gary Houliston
Portfolio Leader – Restoring Ecosystems
Manaaki Whenua – Landcare Research

Contents

Summary iv
1 Introduction 1
2 Background 1
3 Objectives 1
4 Methods and results 2
4.1 Literature database 2
4.2 Point location mapping 2
4.3 Mapping dune ecosystems 3
4.4 Access to naturally uncommon ecosystem maps 6
4.5 Updated factsheet descriptions 7
4.6 A matrix key to the ecosystems 7
5 Discussion 11
5.1 Mapping 11
5.2 The key 12
6 Acknowledgements 13
7 References 13

Summary

Project and client

- Manaaki Whenua – Landcare Research (MWLR), a group of the Bioeconomy Science Institute, and our Otago and Greater Wellington Regional council champions on behalf of the Biodiversity Working Group at Te Uru Kahika, were awarded an Envirolink Tools grant to collate information on, selectively map, and operationalise the definitions of New Zealand's naturally uncommon ecosystems

Objectives

In this project we sought to:

- collate a database of published and unpublished literature on naturally uncommon ecosystems
- identify locations of naturally uncommon ecosystems from the literature and map
- improve on current mapping of dune ecosystems to produce a comprehensive national map
- enable access by councils to the joint Department of Conservation (DOC)/MWLR maps of naturally uncommon ecosystems
- develop updates to the nationally uncommon ecosystem factsheets
- develop a matrix key to the naturally uncommon ecosystems.

Methods and results

- We collated a list of published and unpublished literature containing information about naturally uncommon ecosystems.
- Point locations of naturally uncommon ecosystems were mapped using information contained in the literature.
- Using coastal LiDAR we trained a machine-learning algorithm to identify the coastal dune ecosystem complex. We extrapolated this on a national scale to produce a map of coastal dune ecosystems.
- Access to existing maps jointly developed by DOC and MWLR was secured by negotiating a memorandum of understanding (MOU) between DOC and councils. This MOU is on a council-by-council basis and must be signed individually to become active.
- We updated the naturally uncommon ecosystem factsheets and uploaded the results to the factsheet website.
- A matrix key to identify naturally uncommon ecosystems was developed using LUCID software. The key contains definitions of the terms, images, links to the factsheets, and other information.

Conclusions

- The tools developed in this project have been made available on the web or through Te Ura Kahika's Biodiversity Working Group.

1 Introduction

In 2023 Manaaki Whenua – Landcare Research (MWLR), a group of the Bioeconomy Science Institute, and our regional council champions were awarded an Envirolink Tools grant to collate information on, selectively map, and operationalise the definitions of New Zealand's naturally uncommon ecosystems. The grant was championed by Otago and Greater Wellington regional councils on behalf of the Biodiversity Working Group at Te Uru Kahika. This report outlines the methods and products developed as part of this project.

2 Background

New Zealand's naturally uncommon ecosystems are ecosystems that historically occupied less than 0.5% of New Zealand's land area (<134,000 ha; Williams et al. 2007). They were historically rare because they are defined by combinations of environmental characteristics that are unusual in the landscape. Many of these ecosystems have chemical or climatic limitations on plant growth, which mean they do not support tall forest. Where they occur in proximity to human activities they can also be very vulnerable to clearance or degradation (Holdaway et al. 2012).

Although naturally uncommon ecosystems have been formally defined (Williams et al. 2007), those definitions need to be operationalised to enable non-experts to identify them on the ground. Many are poorly mapped, or not mapped at all, and their intergradation with other ecosystems (often highly invaded vegetation communities) and lack of clarity about some of the terms in the definitions create problems for identification. Knowing what and where they are is a key first step to protect them from further degradation and loss, and thereby enable councils to fulfil their many obligations to biodiversity in their regions.

This report details work undertaken by MWLR to make available published and unpublished resources and maps on naturally uncommon ecosystems for councils, and to operationalise the definitions to provide greater clarity to council staff and environmental consultants for the purpose of improving identification in the field.

3 Objectives

In this project we sought to:

- collate a database of published and unpublished literature on naturally uncommon ecosystems
- construct point maps of the locations of naturally uncommon ecosystems from the literature
- improve on current mapping of dune ecosystems to produce a comprehensive national map
- enable access by councils to the joint DOC/MWLR maps of naturally uncommon ecosystems
- develop updates to the nationally uncommon ecosystem factsheets
- develop a matrix key to the naturally uncommon ecosystems.

4 Methods and results

The purpose of this document is to create a record of what was undertaken for this project and therefore we combine the methods and results (product descriptions and outcomes) in the same section.

4.1 Literature database

We compiled a database of articles, books, reports, and grey literature that report on projects that have investigated naturally uncommon ecosystems.

As a starting point we took the literature identified in the online factsheets,[1] which were developed as part of the of the original research programme on naturally uncommon ecosystems. We then used Web of Science (Clarivate USA) to search for associated articles. A further search used the key words related to their names and synonyms of their names. Many of these investigations predate the formal description of naturally uncommon ecosystems, so it was necessary to search widely using their former descriptive names and synonyms. This database is not an exhaustive list of the research on naturally uncommon ecosystems; rather, it aims to provide information on the breadth and the status of the ecosystems.

All identified references were collated in a table, along with the name of the naturally uncommon ecosystem associated with it. There are 470 items covering 66 of the original 71 naturally uncommon ecosystems. For the remaining five ecosystems (calcareous screes, screes of acidic rocks, ultrabasic screes, boulderfields of quartzose rocks and sinkholes) we didn't find information specifically on them, but they are sometimes obliquely mentioned in discussions about the rock type, e.g. information on ultrabasic rocks might mention screes are a potential landform. On average we found nine references per NUE. We found the most information about inland salt pans (30 references) and Braided riverbeds (28 references). The data are available as an Excel file (NUE mapped reference list.xlsx). Copies are held by the authors at MWLR or may be accessed through Te Ura Kahika's Biodiversity Working Group.

4.2 Point location mapping

We searched the literature identified in section 4.1 to identify mappable coordinates for naturally uncommon ecosystems. Where precise coordinates were given, these were used directly. However, in most cases it was necessary to identify locations from maps and manually digitise them. We relied on the literature to provide a correct identification of the naturally uncommon ecosystem. Locations and identities were not ground-truthed, and so their presence, as mapped, may be historical, dating from the time the study was carried out. We prioritised digitising ecosystems for which locations were less well known. We then worked through the remaining ecosystems as time allowed, with the wetland ecosystems our lowest priority as many of these are being worked on in

[1] https://www.landcareresearch.co.nz/publications/naturally-uncommon-ecosystems

other projects. This meant that six ecosystems received little attention: Domed bogs, Lagoon margins, Lake margins, Pahiki, Estuaries, and Braided riverbeds.

At the time of publication of this report we have searched 385 items (82% of the 470 items identified above) and digitised 2,371 points. ecosystems. The number of points per NUE varies between 1 (Blanket mires) and 593 (Active sand dunes). The number of points reflects not just the number of locations but also how well studied a NUE is. For example, there are a handful of inland saline (salt pans) locations but these locations appear in multiple publications. In a few cases, references were reports detailing mapping projects; in these cases we noted the report in the database but did not re-map the data. The data are available as a GIS point file (NUE_mapped_references.gpkg) or an Excel table (NUE mapped reference list.xlsx). Copies are held by the authors MWLR or may be accessed through Te Ura Kahika's Biodiversity Working Group.

4.3 Mapping dune ecosystems

Dune ecosystems were mapped using a machine-learning algorithm. The dune ecosystems comprise several naturally uncommon ecosystems, including active and stable dunes, dune deflation hollows, damp sand planes, dune slacks, and ephemeral wetlands. Many of these change at scales that are not readily mappable without detailed on-the-ground knowledge and may have fuzzy boundaries, so we chose to map them collectively using high resolution digital elevation data.

4.3.1 Image data

We collated digital elevation data from 11 New Zealand regions: Auckland, Bay of Plenty, Gisborne, Hawke's Bay, Northland, Otago, Southland, Taranaki, Waikato, West Coast, and Wellington. The elevation data were derived from aerial LiDAR surveys following the approach used by Schindler (2024). Ground-classified LiDAR point clouds were rasterised into a digital elevation model (DEM) at 2 m ground resolution. The DEM was then converted into a shaded terrain map (*hillshade*) and a slope map. The three layers for each region were combined to create a three-band raster image as primary input to the sand dune detection model.

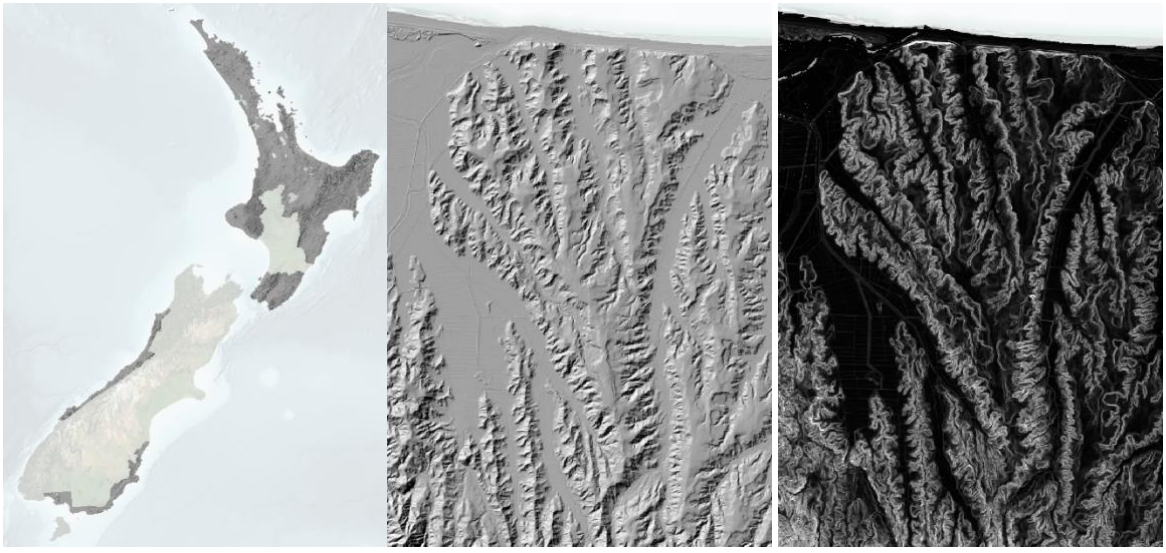


Figure 1. Left: digital elevation data for 11 New Zealand regions; elevation data were then converted into shaded terrain imagery (middle) and slope imagery (right).

4.3.2 Deep learning for semantic segmentation

Deep-learning approaches, especially encoder–decoder convolutional neural networks (CNNs), have quickly become the state of the art in computer vision tasks such as semantic segmentation (e.g. for land-cover mapping in aerial imagery) because of their superior classification accuracy and low level of pixel noise. Deep-learning encoder–decoder networks translate one image to another, with the second image being 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 mapping automation by training the encoder–decoder network on a rasterised version of the current map and the imagery originally used to generate it. The model is trained on a subset of the area of interest where examples 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, in which each pixel represents the class of the corresponding image pixel. However, the model can also generate a probability raster, where intensity represents the model's confidence in each pixel being of the target class. 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 and corrected when the model is run over the same imagery. This approach has the potential to discover additional (i.e. not previously labelled) target objects within the training area.

4.3.3 Model development

We follow a semi-supervised object delineation approach in which we (i) manually label an extensive training data set of sand dunes, then (ii) train a deep-learning model to recognise land-cover class purely from the three-band elevation-derived imagery. The training data were designed to capture a variety of different scenes (local configurations of dune and non-dune landforms) and landscapes, and also areas in which sand dunes are not present (Figure 2).

Manual digitisation is complicated by the fact that many subjective decisions have to be made when drawing polygons. These decisions depend on the local context and many factors interact. In response we have created a training data set that provides a general interpretation of sand dune systems.

The manual digitisation was conducted in a GIS program, and all visible sand dunes within the training areas were delineated. In total, 101 sand dune polygons (with a total area of 191.76 km^2) of different sizes were mapped in 68 training areas, covering a total area of 897.13 km^2. The local regional aerial photography was overlaid on the elevation-derived imagery as a visual aid to guide the manual mapping.

The manual training data generation was done using an iterative process. This workflow involves (i) digitising within one training area, (ii) training an initial deep learning, and (iii) predicting sand dunes over the entire image extent. This process was repeated five times to come to a final model and predictions. We intentionally included many areas such as agriculture and hill country because these were sometimes confused with sand dune systems by the model.

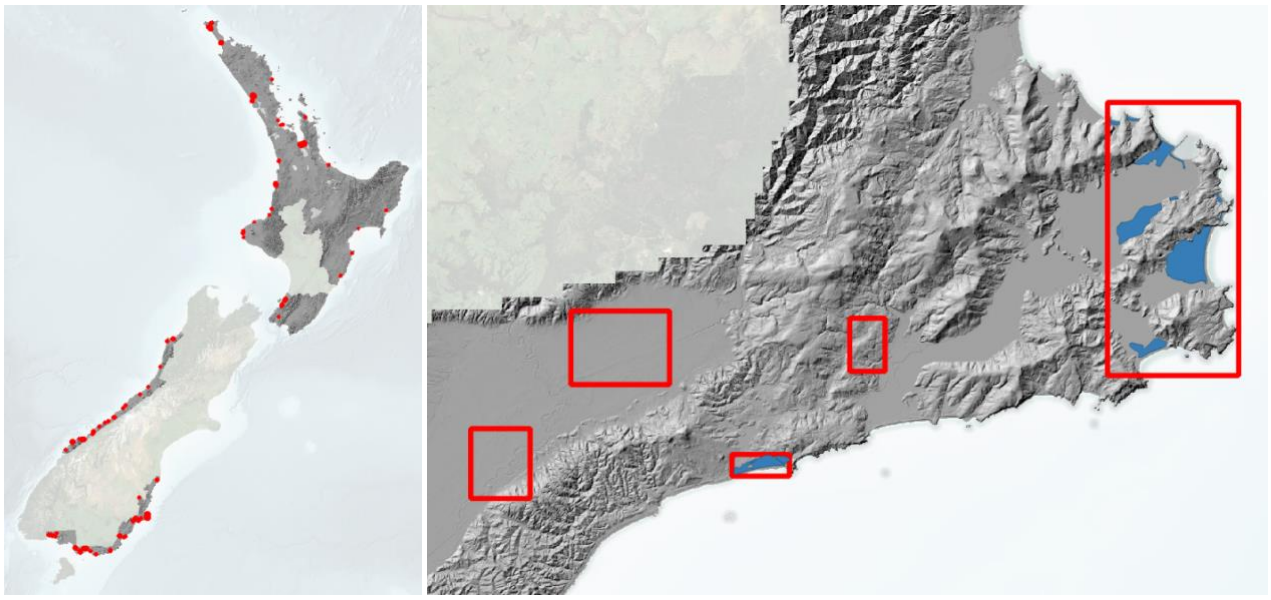


Figure 2. Sand dune training areas (red outlines) and manually digitised sand dune polygons (blue polygons from the Dunedin area).

4.3.4 Training data generation

The three-band raster imagery was split into small square tiles within each training area with 512 × 512 pixels, translating to 1,024 × 1,024 m on the ground (Figure 3). We conducted a random sampling approach, splitting up the image and label data set into 70% training and 30% validation sets. Tiles without valid raster data were excluded from the data sets.

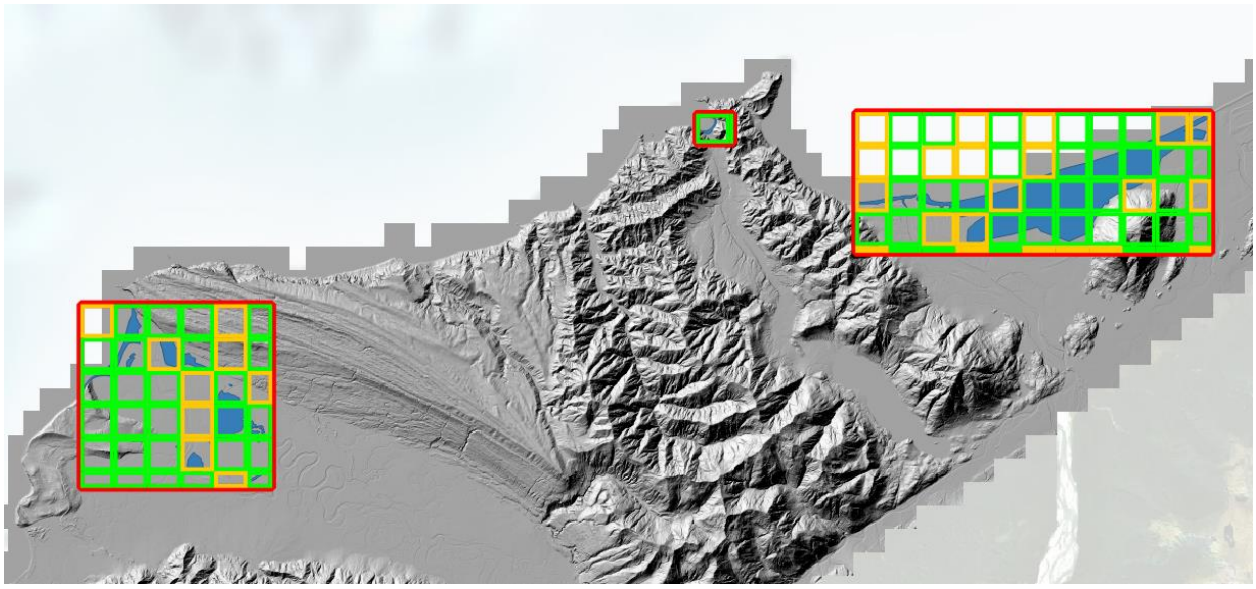


Figure 3. Example of training (green) and validation (yellow) tiles used for the modelling. Red polygons outline the study areas. Tiles without valid raster data were excluded from the data sets.

4.3.5 Model settings

We trained a UNet-type CNN model (Ronneberger et al. 2015) with a Resnet-50 encode backbone pre-trained on the imagenet data set, using the PyTorch framework with focal loss, with gamma = 2, a batch size of 16, a learning rate of 0.0001, and the following augmentations: horizontal and vertical flip, transpose, random rotation (within 90 degrees), grid shuffle (dividing one image tile into nine smaller tiles and randomly shuffling them), and cut out (setting a random square region within a training tile to zero). We trained the model for over 3 hours on an NVIDIA A100 GPU until the validation loss and accuracy reached their maximum and stabilised. We halved the learning rate when the model reached a plateau (i.e. the point at which model loss does not improve by more than 0.00001 after 20 epochs). The final trained model reached a high F1 score of 80.6 %, which measures the segmentation accuracy (i.e. the accuracy of classifying into dune/non-dune) of the validation data set.

When evaluating accuracy metrics such as the F1 score, it is important to recognise that achieving 100% agreement with manual labels is neither feasible nor desirable, because manual digitisation is prone to errors and subjectivity. In some cases model predictions may even outperform human experts, and visual assessments of the model predictions are often more important than pure statistical evaluation.

For inference (i.e. predicting on all imagery for the 11 regions in New Zealand) we applied the trained model to a moving window with image tiles of size 1,024 × 1,024 pixels and an overlap area of 256 pixels. The mapped area of dune ecosystems is 65,214 ha across the North and South Islands, but this excludes the major dune system in Manawatū-Whanganui and dune systems in Nelson/Canterbury because the LiDAR imagery was not readily available at the time of the work. However, an update to include those areas is planned. The map is available as a shape file ("sanddunes.gpkg"). Copies are held by the authors at MWLR or may be accessed through Te Ura Kahika's Biodiversity Working Group.

4.4 Access to naturally uncommon ecosystem maps

Access to the naturally uncommon ecosystem maps jointly developed by MWLR and DOC in 2011/12 was negotiated with DOC, who are the custodians of the data. The resulting memorandum of understanding (MOU) enables councils to use the maps, while accepting full responsibility for their accuracy and the use to which they are put. The MOU needs to be signed individually by councils to come into effect, and then the maps can be obtained directly from DOC. As of August 2025 the DOC contact is Kate McNutt.

During the work it became apparent that many of the larger councils had their own detailed mapping of naturally uncommon ecosystems. These councils were often very generous in sharing this information with us, but we have not included it as part of this Tools project because the information is often confidential. However, for future applications, such as threat classifications, we note that this information exists and that it may be accessible on request from the councils themselves.

4.5 Updated factsheet descriptions

The online factsheets hosted by MWLR[2] were updated to clarify features that would enable better identification in the field. Where appropriate this included clarifying limits on extent, additional information on what ecosystems might share similar features, and the key identification characteristics.

The new information captures additional updates arising from this project and a parallel project (funded through envirolink[3]) that aimed to update the list of ecosystems (Brownstein et al. 2025). The updated factsheets are now live on the existing naturally uncommon ecosystem factsheet website.

4.6 A matrix key to the ecosystems

We developed a matrix key to the naturally uncommon ecosystems. A matrix key enables the user to start with whatever information they have and arrive at an identification, rather than having to follow the strict order and information requirements of a dichotomous (binary choice) key.

The purpose of developing a key was to provide a framework to operationalise the definitions of the ecosystems, such that there is a one-stop-shop for descriptions, definitions, and images of the ecosystems to enable better identification of them in the field. The key is designed to discriminate naturally uncommon ecosystems, but not necessarily differentiate them from common ecosystems or once-common but now rare ecosystems. In some circumstances use of the key could lead to misidentifying an ecosystem as a potential naturally uncommon ecosystem when in fact it is one of these other more common types. Therefore it is essential that the additional information provided in the factsheets and in the key be used to help with discrimination. It is up to the user to confirm that the ecosystem is the candidate ecosystem, based on the resources provided in the key, their own expertise, and any additional online resources they can access.

4.6.1 Coding of features

We began by coding the descriptors for the definitions of the naturally uncommon ecosystems provided by Williams et al. (2007). These were largely accepted as is, although for a small number of ecosystems it became clear, based on updated information, that some modification to the original definitions was necessary so that these ecosystems could be unambiguously identified in the field. We also uncovered the occasional anomaly in the original definitions, which will be addressed in a future revision of the typology.

It soon became clear that it would be necessary to code for all features for each ecosystem because absences of features are the critical diagnostic in the key. This ensures that when a user fails to record a feature, the key will retain the ecosystem because the user may not have assessed it or be able to assess it. Conversely, if a feature is noted that is not a characteristic of the ecosystem, then

[2] https://www.landcareresearch.co.nz/publications/naturally-uncommon-ecosystems

[3] https://www.envirolink.govt.nz/assets/Envirolink/2444-ORC14-Natural-uncommon-ecosystems-recommendations-for-updating-the-list.pdf

the appropriate behaviour of the key should be to omit the ecosystem from the list of candidate identifications. We found it necessary to include additional landforms to cover possibilities that were not present in Williams et al. (2007).

We created one further feature set called 'Geomorphological process'. This feature is a plain-language description of the geomorphological process affecting the surface. While some of this information is often available by interpretation from other features, such as landform and specific parent material, this feature was intended to arrive at an identification faster.

For some features there is a natural hierarchy, such as for the parent materials. For example, calcareous rocks is a higher-level grouping, which can be further split into the specific materials limestone, dolomite, and marble. Where a specific material was part of the definition in Williams et al. (2007), we included the higher-level parent material by default. For example, 'calcareous coastal cliffs' lists the parent material as "limestone rock" (Williams et al. 2007), therefore we also enabled the user to tick 'Calcareous rock" without leading to the exclusion of 'calcareous coastal cliffs'. Conversely, where the high-level parent material only was given, then all specific materials within that group were generally included by default. In some cases we refined this to exclude particular materials, especially when dealing with the 'Other' category. This category contains a diverse assortment of materials, with different origins, and so we were more specific.

For geothermal systems there are several features that are specific to discriminating entities within this group. These features have been incorporated as a folding key so that they become visible if 'geothermal' is selected.

4.6.2 Construction of the key

The completed score matrix was then imported into the statistical package R (R Foundation for Statistical Computing, Vienna, Austria[4]), and reformatted and renamed so that the appropriate hierarchies would be created by the key software. The resulting feature-by-entity matrix was then imported into the LUCID Builder package (Ver 4.0.42; Identic Pty Ltd 2025) as a matrix key.

Definitions for terms used in the key were developed and linked to the terms in the key. The definitions were obtained from a variety of sources, including Williams et al. (2007). We kept to the intent of the definitions in Williams et al. (2007) but interpreted and clarified where needed. Where we had images to illustrate these definitions, these are included.

For each naturally uncommon ecosystem we provide a link to the online factsheet, images that illustrate the ecosystem to help identification, and a partial re-creation of the relevant factsheet for when the user of the key is offline (mini-factsheets).

The key was exported as a stand-alone application with three components:

- "naturally_uncommon_ecosystems_key_v5.html" which is the executable that runs the key
- A folder called "key" which contains all the data for the key including the definitions .html pages and those for the mini-factsheets.

[4] https://www.R-project.org/

- A folder called "player" which contains the LUCID player software to run the key.

The key requires an installation of Java (Oracle Corporation, California) to run. It will be made available on the naturally uncommon ecosystem factsheet home page[5] and the MWLR 'Identification tools' page.[6] The data matrix that underlies the key is available from the authors or through Te Ura Kahika's Biodiversity Working Group ("Data matrix for the key_V5.csv").

The key is being rolled out progressively. The first stage is limited roll out to council staff through Te Ura Kahika's Biodiversity Working Group, including a presentation to this group. After testing by these users we will make it more widely available to the public through the Manaaki Whenua Group's website in the middle of 2026. The reason for this additional delay is primarily to ensure that the revision of the naturally uncommon ecosystem framework is published so that there are not competing versions of ecosystem names, definitions and factsheets in the public arena. This release will be accompanied by a LinkOnLine webinar to reach a wider audience and introduce all the updates at once.

[5] https://www.landcareresearch.co.nz/publications/naturally-uncommon-ecosystems

[6] https://www.landcareresearch.co.nz/tools-and-resources/identification

Table 1. The features coded for each naturally uncommon ecosystem in the LUCID key

Climate regime	Climate local	Geomorphological process	Parent material	Dominant particle size	Soil age	Drainage	Chemical environment	Disturbance regime	Landform
Coastal Inland Alpine	>200 frost days per year Extreme wind exposure Late snow lie Semi-arid (high–very high annual water deficit) None of the above	Coastal sediment deposition and/or erosion Gravity-driven deposition Hydrothermal Running water erosion and/or deposition Shaped or directly deposited by glaciers Volcanism Wind deposition and/or erosion None of the above	**Acidic rocks** Granite and gneiss Mudstone (hard) Mudstone (soft) Rhyolite Sandstone (hard) – greywacke Sandstone (soft) Schist Unknown/other acidic rocks **Basic rocks** Andesite Basalt Diorite Gabbro Tuffaceous mudstone Tuffaceous sandstone Unknown/other basic rocks **Calcareous rocks** Dolomite Limestone Marble Unknown/other calcareous rocks **Quartzose rocks** Quartz sandstone (soft) Quartzite Unknown/other quartzose rocks **Ultrabasic rocks** Ultrabasic Unknown/other ultrabasic rocks **Other or non-specific** Alluvium and till Dune sand Peat Shells Volcanic	Silt and clay (<62.5 µm) Sand (62.5 µm – 2 mm) Gravel (2–64 mm) Cobbles (64–256 mm) Boulders (>256 mm) Bedrock (*in situ*)	Raw Recent Mature Over-mature	Excessive drainage Near permanently saturated (but water table not high) Permanently high water table Regularly high water table Seasonally high water table (periodically dry) Adjacent to open water None of the above	Atmospheric salinity Enhanced nutrients Animal excreta Nutrient-rich groundwater **Geothermal** Acid rain Acid soils, toxic elements Excessive heat Superheated steam Ground water salinity	Periodic flooding (e.g. overflow from stream/river, tidal or wave wash) Seabirds – burrowing Seabirds – guano deposits Seabirds and marine mammals – trampling and grazing None of the above	Beach Beach ridge Cave entrance Cliff Depression Doline Dome Dune (crest and slope) Estuary Fan Gorge Headland Hillcrest Hillslope Lagoon Moraine Plain Plateau Platform Scarp Streamside Subterranean Talus slope Terrace Tor

5 Discussion

5.1 Mapping

All maps in this Tools project, either collated or generated by the project, should be assumed to contain error. The maps are intended to indicate the presence of the ecosystem in the vicinity of the mapped unit, and the boundaries or exact point locations may not be precise. If accuracy is required, such as when evaluating the impacts of development in the vicinity of an ecosystem during consenting, then a site visit is necessary.

Most of the maps are at best only partially truthed. For example, although technically the points were ground-truthed when the associated study was completed, there is no guarantee that the ecosystem will still be there and the process that created it will be functional. The DOC/MWLR maps have, for the most part, not been ground-truthed at all. The dune ecosystem map generated in this project has been partially truthed as a desktop exercise.

We applied two different mapping techniques: point locations based on published and unpublished expert knowledge and machine learning. The point location approach has been excellent for mining existing knowledge, some of which is now dated. Although it doesn't map the boundaries of ecosystems, this is perhaps a more realistic approach given the challenges of doing such mapping accurately. However, users should be aware that the map contains historical information and may not represent the current state. In addition, the literature database is unlikely to be exhaustive for all ecosystems; there are likely some unknown number of additional publications in the grey literature such as DOC internal reports that we did not access.

Machine learning shows a lot of promise as a tool to increase the accuracy and reach of national-scale mapping. Our dune ecosystem maps were comparatively accurate (80.6 % accuracy at correctly classifying dune vs non-dune ecosystems in the validation dataset) because the LiDAR was very successful at capturing the dune landform. Improvements in the model could be achieved by bringing in other predictors such as the underlying geology from QMAP[7]. Extending the predictions to the rest of New Zealand requires a fine scale digital elevation model based on the LiDAR data for the remaining areas (Canterbury, Marlborough, Nelson, Manawatu and Fiordland).

Breaking the dune landform down into its constituent naturally uncommon ecosystems would have been challenging. Additional information would certainly have been required, and it is questionable that this would have been available at resolutions that would allow us to differentiate the ecosystems, given that they grade into one another, often at fine spatial scales. This approach could work on other aerially distinctive entities that have either a distinct geomorphological signature (screes and moraines) or seasonally visual contrasts (e.g. ephemeral wetlands). However, given that naturally uncommon ecosystems are not defined based on vegetation, approaches that characterise plant cover are likely to perform poorly.

Many of these ecosystems grade into others at quite fine scales. This is a challenge for accurately mapping boundaries. An alternative approach is to map the probability of an ecosystem occurring

[7] https://www.gns.cri.nz/data-and-resources/digital-qmap-geological-maps-at-1250000/

at a location. This would encourage site visits when assessing environmental effects of development proposals, while still flagging risks, and be less offensive to some landowners who have concerns about the limits placed upon them by having ecosystems on their properties.

5.2 The key

The key will help differentiate between different naturally uncommon ecosystems, but it is not intended to definitively differentiate natural uncommon ecosystems from other ecosystems. Having come up with an identification, the user should consult the supporting information. These links to images and factsheet information for each of the entities are an important part of the key and should be used to confirm an identification. Uninformed use of the key may lead to incorrect identifications.

Extensive testing was conducted, both in the field and by scoring the features from photos of known locations of field-identified naturally uncommon ecosystems. However, no process can be exhaustive, and there may be combinations of characters that relate to particular ecosystems that we have not thought of or encountered, or ways of interpreting landscapes that we haven't considered. We hope it will be generally useful, but some interpretation may be required on the part of the user in some instances. We welcome feedback that can be used to improve the key.

Determining the boundaries of a naturally uncommon ecosystem can be one of the hardest challenges. Often the boundaries are not sharp and invasions can exacerbate this. At fine spatial scales there can often be several ecosystems mixed together, such as in the dune complexes. As a rule of thumb, the ecosystem stops where the *defining drivers* stop. Hence the following guidance applies.

- Cliffs, tors, scarps and screes: the ecosystem stops at the boundary of the driver. Individual tors are naturally uncommon ecosystems within a matrix, as is the cliff, but not the ground below it, unless it happens to be a talus slope. Notwithstanding this, it may still be desirable to manage the entire landscape containing the ecosystems (including the matrix) to maintain ecosystem integrity.
- Birds and marine mammals use a variety of ecosystems, and therefore a site can be more than one kind of naturally uncommon ecosystem. For example, a cliff might also be a roosting site for a seabird.
- The naturally uncommon ecosystem classification system is terrestrial, so in the case of permanent water bodies the ecosystem applies to the margins (which may themselves be regularly inundated) and not the area permanently occupied by the water body itself. In contrast, the bed of an ephemeral wetland, or a braided riverbed channel that is occasionally dry, is part of the terrestrial ecosystem.
- Ecosystems are not described based on vegetation. In fact, there does not have to be vegetation for it to be a naturally uncommon ecosystem. Obviously, this is the case for subterranean ecosystems, but this is also possible for high-disturbance sites such as shingle beaches or active dunes.
- Judgement will be required in many cases as to identify a boundary, through careful observation of the extent of the influence of the defining ecosystem drivers. For some ecosystems these boundaries will be dynamic, changing over time.

In summary, the naturally uncommon ecosystems were defined based on the drivers, not the vegetation. As a result the ecosystem can still be present even if it is heavily invaded, oversown or otherwise not displaying any form of indigenous dominance. This was a deliberate choice by Williams et al. (2007). Restoration to the naturally uncommon ecosystem is still possible when the processes and drivers of the ecosystem are still present. Therefore, it is up to land managers and council staff to identify the biodiversity, or other values present, before making decisions about an ecosystem's future.

6 Acknowledgements

Thanks to our council champions Scott Jarvie and Roger Uys for useful discussions and guidance throughout the life of the project. We are very grateful to the many people in councils who provided us with information, and to Becky Parmenter and James Arbuckle for digitising. Thanks to Ray Prebble for editing, and to Jo Carpenter for review. We thank Mike Aviss, Simon Litchwark and Dhana Pillai for their naturally uncommon ecosystem expertise in the field to help us operationalise the naturally uncommon ecosystem definitions. Thanks to Kate McNutt, David Burlace, Ellery Mayence, Lisa Brown (DOC) and Kate Dougherty (MWLR) for assistance with the MOU for the DOC/MWLR ecosystem maps. Thanks also to Bill Dyck for guidance with the application and to the councils who supported its development.

7 References

Brownstein G, Monks A, Richardson S, Buxton R, Wiser S 2025. Naturally uncommon ecosystems – recommendations for updating the list. Manaaki Whenua – Landcare Research Contract Report LC4578, prepared for Otago Regional Council.

Holdaway RJ, Wiser S, Williams PA 2012. Assessment of New Zealand's naturally uncommon ecosystems. Conservation Biology 26(4): 619–629.

Ronneberger O, Fischer P, Brox T 2015. U-Net: convolutional networks for biomedical image segmentation. In: Navab N, Hornegger J, Wells W, Frangi A eds. Medical image computing and computer-assisted intervention – MICCAI 2015. Lecture Notes in Computer Science, vol. 9351. Springer, Cham. https://doi.org/10.1007/978-3-319-24574-4_28

Schindler J 2024. LiDAR-derived vegetation layers in Hawke's Bay. Manaaki Whenua – Landcare Research Contract Report LC4470, prepared for Hawke's Bay Regional Council. https://www.hbrc.govt.nz/our-documents/reports-search/details/16437.

Williams PA, Wiser S, Clarkson B, Stanley MC 2007. New Zealand's historically rare terrestrial ecosystems set in a physical and physiognomic framework. New Zealand Journal of Ecology 31(2): 119–128.