

Continuous land cover change detection in a critically endangered shrubland ecosystem using neural networks

Glenn R. Moncrieff ^{1,2,*} 

¹ Fynbos Node, South African Environmental Observation Network, Private Bag X7, Rhodes Drive, Claremont 7735, South Africa

² Centre for Statistics in Ecology, Environment and Conservation, Department of Statistical Sciences, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa

* Correspondence: glenn@saeon.ac.za

Abstract: Existing efforts to continuously land cover change using satellite image time-series have mostly focused on forested ecosystems in the tropics and northern hemisphere. The notable difference in reflectance that occurs following deforestation allows land cover change to be detected with relative accuracy. Less progress has been made in detecting change in low productivity, disturbance-prone vegetation such as grasslands and shrublands, where natural dynamics can be difficult to distinguish from habitat loss. Renosterveld is a hyperdiverse, critically endangered shrubland ecosystem in South Africa with less than 5–10% of its original extent remaining in small, highly fragmented patches. I demonstrate that direct classification of satellite image time series using neural networks can accurately detect the transformation of Renosterveld within a few days of its occurrence, and that trained models are suitable for operational continuous monitoring. A dataset of precisely dated vegetation change events between 2016 and 2021 was obtained from daily, high resolution Planet labs satellite data. This dataset was then used to train and evaluate 1D convolutional neural networks and Transformers to continuously detect land cover change events in multivariate time-series of vegetation activity of Sentinel 2 satellite data as it becomes available. The best model correctly identified 89% of land cover change events at the pixel-level, achieving a f-score of 0.93, a 79% improvement over the f-score of 0.52 achieved using a method designed for forested ecosystems based on trend analysis. Models have been deployed to operational use and are producing updated detections of habitat loss every 10 days. There is great potential for supervised approaches to continuous monitoring of habitat loss in ecosystems with complex natural dynamics. A key limiting step is the development of accurately dated datasets of land cover change events with which to train machine learning classifiers.

Keywords: Land cover change; Land cover monitoring; Deep learning; Renosterveld; Threatened ecosystems; Sentinel 2; Planet labs

Citation: Moncrieff, G. R. Continuous land cover change detection in a critically endangered shrubland ecosystem using neural networks. *Remote Sens.* **2022**, *1*, 0. <https://doi.org/>

Received:

Accepted:

Published:

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2022 by the authors. Submitted to *Remote Sens.* for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Land cover change is the single largest cause underlying the global biodiversity crisis [1]. It is estimated that at least 21% of all terrestrial plant extinctions are directly related to agriculture alone [2]. A core component of efforts to halt biodiversity loss and climate change are remote sensing-based monitoring systems built upon satellite-derived earth observation data [3–5]. These systems allow up-to-date reporting on trends in habitat loss, and in some instances, **near-real time alerts** that may aid interventions and enforcement. Typically, temperate and tropical forests are the focal ecosystems for applications, examples of which include **Global Forest Watch** and **Terra-i**.

While forests are indeed globally important stores of carbon and home to myriad species, far less attention has been paid to developing remote sensing-based monitoring for use in open ecosystems where trees may be present but not dominant, such as savannas, grasslands, shrublands and deserts. Open ecosystems cover a large proportion of the globe, make up 40% of the global total ecosystem organic carbon and harbour a substantial

proportion of the world's biological diversity [6,7]). Failure to develop capacity for change detection in open ecosystems would exclude the majority of the earth's terrestrial surface from monitoring. Detecting abnormal change in forested ecosystems is relatively simple in comparison to open ecosystems. The spectral signature of a closed-canopy forest stands in stark contrast to the sparsely vegetated or bare ground that remains after deforestation. In contrast, the natural vegetation state in open ecosystems can vary dramatically due to natural disturbances, long-term trends or cyclical functions, such as those relating to fire, seasonality and interannual weather variations [8,9]. This variability found within healthy ecosystems can obscure spectral change that occurs when intact natural vegetation is converted to non-natural land use.

The Cape Floristic region (CFR) in South Africa is home to some of the world's most floristically diverse open ecosystems [10,11]. Within the CFR the most threatened vegetation type is the grassy shrubland known as Renosterveld [12,13]. Renosterveld is restricted to relatively nutrient-rich soils in the south and west of the Western Cape Province. It is noteworthy for its exceptionally high diversity of bulbs and - due to widespread conversion of intact Renosterveld to agriculture - perhaps the highest concentration of plants threatened with extinction in the world [13–16]. Only 5–10% its original extent remains, with much of this contained in small isolated fragments [17]. Despite legislation prohibiting the destruction of Renosterveld habitat, loss is ongoing [18]. It is estimated that at current rates Renosterveld will become extinct as an ecosystem within the 21st century [13]. The primary cause of this decline is the expansion of intensive agriculture (primarily for the cultivation of grains and oil-seed such as wheat, barely and canola), and overgrazing by livestock.

A remote sensing-based monitoring system, such as is available for many forest ecosystems worldwide, may help to reduce rates of Renosterveld loss by aiding in the identification of land owners responsible for historical loss. If information is provided with low latency between the onset of ploughing and its detection, there may be potential to intervene while habitat destruction is underway. Continuous change detection approaches attempt to assimilate and analyze new remote sensing data as they become available, and have the potential to provide near-real time alerts on land cover change. A typical pattern used for continuous change detection involves analyzing trends in historical vegetation activity and building a model that can accurately forecast the expected reflectance of intact forests. Predictions are then compared recently acquired observations [19–22]. When observed vegetation activity falls outside the range expected for natural vegetation, this is taken to be indicative of forest loss. Trend analysis methods used to model expected vegetation activity typically assume a relatively predictable progression of vegetation activity in intact ecosystems that can be well described by linear models with harmonic terms to capture seasonality [20]. The resulting land cover change detections, while impressive for stable ecosystems such as forest, are prone to commission errors in dynamic ecosystems where natural variability can easily be mistaken for land cover change and omission errors when disturbance is subtle. These problems are particularly prevalent in non-forest ecosystems with natural dynamics such as fire, aridity and herbivory [8].

When ground-truth data that distinguish areas where habitat loss has occurred are available, land cover change can be directly classified from time-series of vegetation activity using supervised learning. This removes the need to specify a model for expected vegetation activity [23,24]. For example, Hansen *et al.* [25] trained decision-tree classifiers to detect forest disturbance in the tropics, comparing recent imagery to a stable reference period. Their system is widely used for forest monitoring in the tropics and has been incredibly impactful. These approaches are however limited to producing updates on land cover change at discrete intervals (often annually), as the training data typically do not specify the exact date on which change occurred, but rather that change occurred at some unknown point between two fixed points in time (e.g. January and December). While this provides useful information for reporting on habitat loss, this information is usually only obtained months or years after the reported change has occurred [24,26].

If a precise date for a land cover change event is available, supervised methods can be trained to classify change using input data covering a period of time shorter than year (Figure 2). Furthermore, if these labelled events occur at a variety of points in time throughout the year (and potentially across multiple years) models that can recognise the signal of land cover change regardless of the time of year and are robust to intra and inter-annual natural variability can be trained [27]. These models can then be applied continuously as new data become available and produce near-realtime land cover change alerts. They would also be less prone to errors than change detection methods requiring the ecosystem to follow a predictable pattern (e.g Verbesselt *et al.* [19], Zhu and Woodcock [20], Zhu *et al.* [22]), as the models could theoretically learn to distinguish the spectral patterns indicative of land cover change from natural disturbances at any point in time if shown sufficient examples.

Machine learning approaches such as random forests and support vector machines have been widely used to directly classify land cover change using multi-temporal remotely sensed input data [24,28–30]. These methods do not however explicitly take into account the temporal structure within the data. Neural networks specifically designed for classification of sequential data such as convolutional neural networks (CNN) or Transformers have been shown to improve performance on time series classification tasks relative to other machine learning methods [31]. They have recently been being applied to land cover classification with great success [32–34] and direct classification of land cover change [35].

Here I describe the design and performance of a system for continuous land cover change detection in a critically endangered shrubland ecosystem, the Renosterveld of South Africa. Using a dataset of precisely dated land cover change events covering multiple years, two different deep learning architectures are trained and evaluated against tree-based and trend analysis methods commonly used for land cover change detection and monitoring in forested ecosystems. Model building and design choices are motivated by an intention to produce a system that can be deployed for operational use by conservation management and law enforcement agencies. This demands a system that is robust to the natural variability within this ecosystem, can assimilate and make predictions on new data as it becomes available, and is able to detect habitat loss with minimal delay.

2. Materials and Methods

2.1. Study region

The study region is limited to the extent of the Renosterveld bioregion in the Overberg district municipality in the Western Cape Province of South Africa (Figure 1)[36]. The choice of region was determined by the critically endangered status of Renosterveld vegetation, the availability of accurate land cover data for the Overberg municipality, and partner conservation organizations in the region. The Overberg municipality covers an area of 12,241 km², with Renosterveld covering 4711 km² with this region naturally. Detailed mapping of remaining natural vegetation in the Overberg based on 2003 data identified 656 km² of the original extent of Renosterveld vegetation in the Overberg remaining.

2.2. Land cover change events

Using random forests for change detection on Sentinel 2 imagery and subsequent detailed verification with very high resolution imagery, Moncrieff [18] identified all conversion of Renosterveld to non-natural land cover between January 2016 and January 2020. Land cover change was identified within 268 parcels (a spatially contiguous land cover change event) totalling 478.6 ha. Using daily Planet Labs PlanetScope data, dates were assigned to each event indicating the closest possible timing for the occurrence of the land cover change event. It is not possible to assign every event to a single day as often multiple days or even weeks elapse between the beginning and end of a land cover change event. Furthermore, clouds and haze can reduce the frequency at which clear observations are available. Thus for each land cover change event two dates were assigned - the latest date on which the presence of natural vegetation could be confirmed, and the earliest date on

which the absence of natural vegetation could be confirmed (Figure 2). These two dates occur within one week for 21% of all events, and within two months for 87%. Here a sample of these data are used where land cover changes dates were assigned with the highest confidence, leaving 185 events and 219 ha of land cover change.

This dataset contains only events in which complete conversion of intact natural vegetation occurred and transformation into non-natural land cover could be dated within a short period of time. It thus excludes gradual degradation occurring over multiple months or years. Fire is a natural process in Renosterveld and is not considered a land cover change event unless it precedes further modification and subsequent conversion to croplands. An additional 4 events covering 7 ha were identified within the same region between Jan 2020 and Jan 2021. These events were used for testing model performance outside of the temporal range of training data, providing a more accurate assessment of operational performance. To augment data with additional observations of stable natural vegetation, parcels of vegetation that remained natural over the entire study period were added. This added an additional 2590 ha dataset covering the period 2016–2019 and 24 ha to the testing dataset covering 2020. These no-change parcels were selected from the [17] map of remaining Renosterveld, buffered to ensure that they did not occur within 100m of non-natural vegetation and manually verified.

2.3. Sentinel 2 time series

Input time series' for classification were created using Sentinel 2 L1C data (S2). For training and validation all available S2 data up until 2019-12-31 were extracted and pre-processed, whereas testing data covered the period 2020-01-01 to 2020-12-31. S2 images are pre-screened and discarded if total scene cloud cover exceeds 50 percent. Pixel-level cloud masking is performed using the [s2cloudless](#) machine learning-based cloud masks. Pixels with a cloud probability greater than 40 percent are masked. Cloud shadows are also masked and are calculated from cloud masks using cloud and solar geometry combined with dark pixel detection. All S2 bands are retained and the following additional vegetation indices calculated: Normalized Difference Vegetation index (NDVI), Enhanced Vegetation Index (EVI), Normalized Difference Water index (NDWI), Normalized Burn index (NBR) and Normalized Difference Red-edge index (NDRE). Data are exported at 10m resolution with 20m and 60m bands resampled using bilinear interpolation.

S2 data area available every 10 days in this region until mid-2017 when data becomes available at least every 5 days subsequent to the launch of Sentinel 2B. To create regular time series from the S2 image stacks pixel-level time series are created with a single observation every 10 days. Where multiple images are available within a 10 day period, data from the image with the highest pixel-level NDVI value are used. Where no data is available, values from the previous valid image is used unaltered. In this application interpolation between observations by averaging is undesirable as it may obfuscate sudden changes in reflectance diagnostic of land cover change. Classification is performed on time series of reflectance collected over 180 days, equalling 18 observations with a 10 day interval. The choice of 180 days was arrived upon through experimentation with both longer and short time series'. While it is theoretically possible to detect Renosterveld loss with as few as 2 sequential observations if the change occurs rapidly, few of the labelled events have such precision in determining the actual data on which habitats loss occurred. Longer time series contain more information on the spectral properties of vegetation before and after change, but no additional performance was gained when experimenting with time series longer than 180 days.

2.4. Data labelling

Classification is performed at the pixel-level. The S2 time series for each pixel within a parcel is assigned the label 'natural' if the last date of confirmed natural vegetation falls after the end date of the time series under consideration (Figure 2), or if no change was ever reported within that parcel. If both the last date of confirmed natural vegetation and

the first date of confirmed land cover change occur within the 180 day time series' the land cover change event must have occurred within the time series under consideration and the 'recently transformed' label is assigned. If the last date of confirmed natural vegetation occurs within the time series and the first date of confirmed land cover change falls after the end of the time series it is unclear whether land cover change has indeed taken place and the data point is discarded. Similarly, if the first date of confirmed land cover change occurs before the start of the time series, land cover change has already taken place. Here a stable non-natural land cover class is not included, thus these data are discarded.

The resulting labelled dataset is highly imbalanced, with the vast majority of sampled pixels in the 'natural' class. There is also very high spatial auto-correlation among neighbouring pixels within the same parcel. In addition, successive time series for the same pixel obtained 10 days apart will share 94 percent of observations. To ameliorate class imbalance and spatial autocorrelation only 5 percent of pixels within parcels where no change occurs are sampled and 33 percent of pixels from parcels where change does occur are used. To reduce temporal oversampling within a pixel, time series windows of 180 days are shifted forward by 60 days for every sample taken. Subsequent to this data thinning, pixels are stratified by parcel and assigned to training and validation sets using a 70/30 split - that is 70 percent of parcels are allocated to the validation set and 30 percent to the validation set. Test data is intended to evaluate model performance in an operational setting and hence does not undergo temporal or spatial thinning, with a new prediction made every 10 days.

2.5. Models

Two neural network architectures that have shown excellent performance in classification of land cover from satellite image times series are trained and evaluated along with random forests, a tree based method commonly used for land cover classification. TempCNN is a 1D convolutional neural network for land cover classification from satellite image times series [34]. Convolutional kernels that extract relevant features in local temporal regions across multivariate time series of individual pixels are learned. These 1D CNN layers are followed by non-linear activations and potentially also batch normalization and stacked, before class probabilities are produced by a dense and softmax layer. TempCNN performs well on a range of land cover classifications tasks, and has been adopted by multiple packages for classification of satellite image time series [37,38].

Transformer models were originally proposed as sequence-to-sequence encoder-decoder models based on the self-attention operation and applied to language translation [39]. Attention mechanisms operate by assigning an importance-score (attention) to weight each element in a sequence [40]. In self-attention, the concept of attention is used to encode sequences to a hidden representation. In the traditional application of transformers to language tasks, self-attention encoders are followed by self-attention decoders. For time series classification only the encoder network is needed. The encoder network learns which elements in a time-series are relevant for the classification task at hand and extracts salient features. In this work, a Transformer model that has been evaluated in Rußwurm and Körner [33] is used, where individual pixels' time series' are encoded using stacked self-attention layers. This architecture outperforms many other neural networks and tree-based methods for crop type mapping using satellite image time series [37].

2.6. Model fitting

Hyperparameter optimization for random forests, TempCNN and the Transformer network is performed using a grid search over 24, 36 and 32 parameter combinations respectively. The best fitting model of each class is then evaluated against test data. For the Transformer and TempCNN, the Adam optimizer is used to minimize focal loss. TempCNN and Transformers as implemented in Tensorflow 2 and the eo-flow python package. scikit-learn is used for Random Forests. Google Earth Engine is used for preprocessing of input data [41].

To aid model interpretation, the most important temporal regions identified by the TempCNN model for detecting land cover change are calculated using Grad-CAM++ [42]. Originally developed for use with CNNs and image classification, Grad-CAM++ can easily be used when CNNs are applied to time series classification. Grad-CAM++ calculates a saliency score (importance) for each time step in a time series toward the final prediction made by the model. This is calculated using a weighted average of all positive activation gradients across channels in the final convolutional layer.

2.7. Trend analysis comparison

The accurately dated land cover change data used here to train supervised methods is seldom available. Trend analysis approaches to land cover change detection are more generally applicable and widely used [43,44]. A method based on trend analysis is used to evaluate the improvement in accuracy achieved when using direct classification on these data. The approach followed is very similar to the approach outlined by Zhu and Woodcock [20], termed Continuous Change Detection and Classification (CCDC). CCDC is widely used for continuous change detection, and implemented in platforms such as Google Earth Engine. More recent methods based on similar principles that would improve on this baseline are available [21,45,46], but they are not commonly used on S2 data.

The CCDC approach operates by fitting a simple linear model with a single harmonic term to capture seasonal variation using robust regression through iteratively reweighted least squares with the **Talwar cost function**. The model is fitted to historical NDVI data for every pixel in regions known to still contain Renosterveld. The expected NDVI over the test period (2020-01-01 - 2020-12-31) is then calculated for each day of the year at each pixel. This is compared to the observed NDVI and anomalies calculated and expressed in terms of the RMSE of the fitted regression. If 3 successive observations differ from the prediction by 3 times the RMSE, land cover change is assumed to have occurred. The RMSE thresholds used to trigger detection can be changed in order to account for ecosystems with greater inherent variability [20]. Different thresholds were experimented with, with 3 times the RMSE providing the best performance here. More sophisticated implementations of this workflow make use of multiple bands, or use parameters learned from labelled land cover change events [46].

3. Results

Full time series' of selected pixels extracted for the entire period covered by the training and validation data are shown in **Figure 3**. These data show the highly dynamic and variable nature of vegetation in this ecosystem. In natural vegetation NDVI peaks in winter, corresponding to the wet season, but the onset of greening is highly variable. Summer minimum NDVI is very low, though both minimum and maximums are highly variable across the region. Nonetheless a distinctive signal attributable to land cover change is visible, and models are able to use this signal to reliably detect land cover change.

The best TempCNN model used a dropout rate of 0.8, L2 kernel regularization with a regularization factor of $1e-8$, He normal kernel initialization, 6 stacked convolutional layers each with 32 filters and 256 neurons in the final fully connected layer. The best performing Transformer used a dropout rate of 0.5, 4 layers in the encoder with 4 attention heads, a model depth/query size of 128, no layer normalization and 256 neurons in the encoder feed-forward layer. The best performing random forest model used 100 estimators, and required a minimum of 4 samples to split an internal node.

The TempCNN achieved the highest performance, with a recall of 0.89 and precision of 0.96 when identifying recently transformed Renosterveld and overall f-score of 0.93 (Table 1). The Transformer was only slightly worse than the TempCNN with a recall of 0.88 and precision of 0.97 on the recently transformed class and overall f-score of 0.92, while the Random Forest was significantly worse than both TempCNN and the Transformer, with recall and precision of 0.79 and 0.98 on recent transformation, and overall f-score of 0.88. All models using direct classification outperformed the trend analysis method, which was




Figure 1. The natural extent of lowlands Renosterveld in study region (left panel) and the remaining natural and degraded fragments (right panel)

far more likely to produce both false negatives and false positives with recall of 0.49 and precision of 0.54 on the recently transformed class and overall f-score of 0.52.

Performance for random forests, TempCNN and the Transformer declines when their ability to detect recent transformation within 20 days of change is measured. Again the TempCNN performed best, with recall of 0.64, followed by the Transformer with 0.60 and the random forest with 0.50. Trend analysis classification is not possible within 20 days of change, as 3 anomalous data points are required before change is flagged. This equates to 30 days with the 10 day interval used here.

Figure 4 provides a visual depiction of the time series features characteristic of Renosterveld loss events, determined using Grad-CAM++ on the trained TempCNN. For a single Renosterveld loss event in a pixel viewed from multiple points in time after the event, the same temporal features are focused upon. Across multiple different land cover change events, the network has learnt to consistently focus on time periods in which a drop in vegetation greenness has occurred. Other more nuanced changes in spectral reflectance may also be contributing to the importance of these temporal regions, as the TempCNN is using all input bands and indices to classify change, though only NDVI is depicted here.

A prototype application performing up-to-date inference over the entire study region every 10 days using TempCNN is available at <https://glennwithtwons.users.earthengine.app/view/renosterveld-watch> and an example sequence of predictions over 2 months in an area through the course of a land cover change event is depicted in Figure 5. Inference over the entire remaining distribution of Renosterveld in the Overberg region for a single date reveals multiple as-yet unreported areas where Renosterveld has been transformed to cropland within the 6 month time series of satellite data under consideration. Because models are trained to only make predictions in regions where land cover is intact Renosterveld at the beginning of the time series under consideration, predictions are not reliable when made on other vegetation types, or where Renosterveld has already been converted to agriculture. This is not an issue when up-to-date land cover maps are available. Multiple observed cases of incorrect predictions of Renosterveld loss where land cover has in reality remained as agriculture throughout the time series under consideration highlight the need for updating of existing regional land cover maps.

4. Discussion

Shrublands are globally important repositories of biodiversity and play an important role in regulating planetary biochemical cycles [6,7]. In many regions shrublands are being lost at a rate that exceeds other ecosystems such as forests [47]. Despite this, few studies have attempted to develop algorithms designed specifically for operational monitoring of shrubland loss, ostensibly because many of these ecosystem display complex natural dynamics and have low overall productivity. These characteristics make the methods typically applied in operational forest monitoring perform poorly. This study demonstrates




Figure 2. The approach used for labelling time series to enable continuous classification of change. Each step in the sequence of 19 observations represents an observation obtained for a single pixel at a regular interval (e.g. 10 days). The land cover label for the pixel (identified by the red point in the images) can either be natural, unknown or non-natural. This label is obtained through inspection of time series’ of Planet Labs PlanetScope satellite imagery. For the example pixel indicated here, an image obtained at time step 3 indicates that the land cover is natural on this date, the image at time step 9 indicates the last date on which natural land cover is observed. The first time at which non-natural land cover is observed is step 12. Hence, it is not known whether the land cover was natural or non-natural at timestep 10 and 11. Using a window length of 7 observations and forward step of 4, 4 labelled sequences can be extracted from the full time series. Each sequence is assigned the label of the most recent observation. Provided enough labelled events occurring at different points in time across multiple years, models that can recognise the signal of land cover change regardless of the time of year and are robust to intra and inter-annual natural variability can be trained. These models can then be applied continuously as new data become available to produce near-realtime land cover change alerts.

Table 1. Model performance evaluated on test data. Recall-1-short reports model recall when change has occurred within the first 2 time-steps (20 days), of the 18 step (180 days) time series. Performance of the top-performing model for each metric is highlighted in bold.

	Accuracy	F1	Recall-1	Recall-0	Precision-1	Precision-0	Recall-1-sho
Random Forest	0.98	0.88	0.79	0.99	0.98	0.98	0.50
TempCNN	0.99	0.93	0.89	0.99	0.96	0.99	0.64
Transformer	0.99	0.92	0.88	0.99	0.97	0.99	0.60
CCDC	0.90	0.52	0.49	0.95	0.54	0.94	-




Figure 3. Observed NDVI time series for the full period covered by training and validation data for randomly selected pixels from across the study region. Vegetation in all time series' begins as natural Renosterveld. In panel a, a land cover change event occurs in the period highlighted in red, and thereafter the land cover is non-natural. In b, vegetation remains natural for the entire time series




Figure 4. Saliency of temporal regions in time series for pixels in which land cover change is predicted to have occurred, determined using Grad-CAM++. Higher saliency scores indicates that the region is more influential in determining that land cover change has occurred. Blue lines show NDVI, though saliency is calculated using all input features. In a) a single pixel is shown at four points in time, with each subsequent panel shifted forward 1 step (10 days). In b) time series' from four spatially and temporal independent pixels are shown.




Figure 5. Example sequence of predictions through the course of a land cover change event in Renosterveld shrublands of the Overberg region, South Africa. Upper panels show the predicted probability of land cover change, where predictions are only made in pixels known to have intact natural vegetation at the start of the monitoring period. The lower panels show the most recent Planet labs PlanetScope imagery collected prior to, or on the date for which predictions are made. The detection of Renosterveld loss along the eastern edge of the monitored area is apparent. This detection occurs within 10–20 days of the actual land cover change event.

that reliable detection of land cover change in shrublands is possible, and that detection is possible within days of land cover change events, thus facilitating the use of these algorithms for operational monitoring and enforcement. The key component of this is the continuous use of direct classification, enabled by the availability of hand-labelled change events with the date of each event precisely determined.

Direct classification is rarely used for continuous detection of land cover change. The labels needed to train models require detailed examination of very-high resolution satellite or aerial imagery and dense time series. Very-high resolution data may not be available for the required region, time period or frequency, and the costs associated with these will be prohibitive for many potential use cases. Even when all the necessary data is available, the data labeller must have intimate knowledge of the ecosystem in question in order to determine whether an observed event is indeed land cover change. For example, Renosterveld is a naturally fire-prone ecosystem and burnt vegetation is not necessarily indicative of land cover change. These limitations should not, however, preclude the development of labelled datasets. The availability and cost of very-high resolution and high temporal frequency imagery is improving dramatically, and new programmes are broadening access to these data (e.g. the Norway International Climate and Forest Initiative - Planet partnership). Advances in deep learning performance on land cover classification when labelled data are limited will further reduce data requirements [48,49]. The investment of time required to acquire labels is indeed significant, but for critically endangered ecosystems that are predicted to become extinct within decades based on current trajectories, surely this investment is justifiable. Particularly if there is reason to believe that an operational land cover change monitoring system will slow or reverse these trends.

The volume of labels required to achieve a performant model in the example presented here is surprisingly small - just 219 ha of dated land cover change events. This is sufficient to produce large improvements in performance in deep learning-based classifiers relative to random forests which typically require less data to train. While it is highly unlikely that a supervised approach will be applicable in any region beyond the domain of the training data (thus requiring further painstaking data collection and labelling), meta-learning and transfer learning approaches ought to reduce data requirements for application to new disturbance types and new geographies [50]. Given multiple labelled datasets of dated land cover change from different geographies or ecosystems, models can be trained such that they are initialized optimally and learn rapidly when confronted with a new regime. Particularly promising is model-agnostic meta-learning, which does not necessitate custom designed models for this task, but rather can make use of widely adopted models such as TempCNN or Transformers [48,51].

The models developed here are not only limited in their spatial domain of applicability, but also the types of land cover change events which they are capable of detecting. Only a single land cover change can be identified - the sudden and destructive conversion of natural Renosterveld to non-natural land cover. Other types of change can occur in this landscape, such as gradual degradation through overgrazing or spread of invasive alien plants. It is possible that the same approach could be used to diagnose these types of changes. This would however require a significant increase in the amount of training data and longer input time series, as these types of change occur over periods longer than the 6 month windows input into classifiers. The choice to focus on a single type of rapidly occurring change allowed relatively high performance to be achieved with limited training data and is focused on the actions considered most severe by law enforcement in the region.

A major improvement achieved through the approach presented here is the ability to detect change within a few days of the actual event. Rapid identification is critical in many intended applications of near-realtime continuous land cover change monitoring - such as law enforcement and management intervention. Other approaches typically require multiple consecutive observations confirming the change has occurred before a positive identification is made [20,21]. Generally any land cover change will be complete by the

time these detections are reported. Here change is identified with relative skill with only 2 observations subsequent to the actual change event. In certain applications this may be sufficient to lead to an on-the-ground intervention that could preventing further damage. Accuracy improves as further observations are acquired, and hence it is possible to delay the identification of change depending on the cost/benefit of any particular application.

The improvement in performance obtained when using deep learning-based time-series classifiers vs. random forests is likely a result of their ability to retain temporal information and extract features based thereon. Previous work comparing the performance of Transformers and TempCNN for land cover classification found Transformers to generally perform better than TempCNN, particularly when minimal preprocessing and no atmospheric correction is performed [33,37]. Here, extensive feature engineering has been conducted and cloud masking performed to remove noisy data, potentially resulting in the convergence in performance between TempCNN and the Transformer. A range of machine learning methods that are not based on neural networks have been developed for the classification of time-series data, and there may be potential performance gains when using these methods, particularly when training data is scarce [52,53]. Gains in performance cannot, however, come at the costs of increases in model inference time to the extent that prediction cannot be made at the required latency or within available budgets. This is particularly relevant because monitoring will often be required over very large regions. Using the TempCNN model fitted here prediction over the entire remaining range of Renosterveld in the Overberg region can be achieved within 120 CPU hours. The full prediction pipeline for preprocessing, prediction and postprocessing is achieved at minimal cost using tools such as Google Earth Engine and Google Cloud Dataflow.

5. Conclusions

The potential of continuous change detection using deep learning on satellite image time series, particularly in non-forest ecosystems, remains largely unexploited. This study demonstrates that when labelled data are available, accurate and precise detection of land cover change is possible at low latency. Obtaining the labelled data does indeed require significant effort, and it is likely that independent datasets and models for will be required for different ecosystems. This should not discourage the further development and operationalization of this approach. Both deep learning and the explosion of remote sensing data have been lauded as having great potential to address the biodiversity and climate crises. For these technologies to truly have impact they must find application beyond the forested ecosystems of the world, and embrace the complexity of open ecosystems.

Funding: This research was supported in part by by the National Research Foundation of South Africa through (Grant No. 118593) as part of the RReTool: Rapid and repeatable tools for monitoring and mitigating global change impacts on natural resources project, the Group on Earth Observations-Google Earth Engine Programme, and the NASA Ecological Forecasting Team Applied Sciences Program (80NSSC21K1183).

Data Availability Statement: Code and data used in this analysis are available at <https://github.com/GMancini/monitor>. Due to ongoing criminal investigations location data has been removed to protect the identity of landowners where necessary.

Acknowledgments: I am very grateful for the support and encouragement given by Odette Curtis and her passion for conservation in this region. Marcel Gietzmann-Sanders helped produce the prototype operational monitoring application. Multiple individuals from the Western Cape Department of Environmental Affairs and Development Planning are acknowledged for their input, in particular, those from the Environmental Law Enforcement directorate.

Conflicts of Interest: The author declares no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

References

1. Brondizio, E.S.; Settele, J.; Díaz, S.; Ngo, H.T. Global assessment report on biodiversity and ecosystem services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services **2019**. 439
2. Le Roux, J.J.; Hui, C.; Castillo, M.L.; Iriondo, J.M.; Keet, J.H.; Khapugin, A.A.; Médail, F.; Rejmánek, M.; Theron, G.; Yannelli, F.A.; et al. Recent Anthropogenic Plant Extinctions Differ in Biodiversity Hotspots and Coldspots. *Current Biology* **2019**, 29, 2912–2918.e2. doi:10.1016/j.cub.2019.07.063. 440
3. Turner, W.; Spector, S.; Gardiner, N.; Fladeland, M.; Sterling, E.; Steininger, M. Remote sensing for biodiversity science and conservation. *Trends in ecology & evolution* **2003**, 18, 306–314. 441
4. Yang, J.; Gong, P.; Fu, R.; Zhang, M.; Chen, J.; Liang, S.; Xu, B.; Shi, J.; Dickinson, R. The role of satellite remote sensing in climate change studies. *Nature climate change* **2013**, 3, 875–883. 442
5. Turner, W. Sensing biodiversity. *Science* **2014**, 346, 301–302. 443
6. Bond, W.J. *Open ecosystems: ecology and evolution beyond the forest edge*; Oxford University Press, 2019. 444
7. Poulter, B.; Frank, D.; Ciais, P.; Myneni, R.B.; Andela, N.; Bi, J.; Broquet, G.; Canadell, J.G.; Chevallier, F.; Liu, Y.Y.; et al. Contribution of semi-arid ecosystems to interannual variability of the global carbon cycle. *Nature* **2014**, 509, 600–603. doi:10.1038/nature13376. 445
8. Slingsby, J.A.; Moncrieff, G.R.; Wilson, A.M. Near-real time forecasting and change detection for an open ecosystem with complex natural dynamics. *ISPRS Journal of Photogrammetry and Remote Sensing* **2020**, 166, 15–25. doi:10.1016/j.isprsjprs.2020.05.017. 446
9. Wilson, A.M.; Latimer, A.M.; Silander, J.A.J. Climatic controls on ecosystem resilience: Postfire regeneration in the Cape Floristic Region of South Africa. *Proceedings of the National Academy of Sciences of the United States of America* **2015**, 112, 9058–9063. doi:10.1073/pnas.1416710112. 447
10. Myers, N.; Mittermeier, R.A.; Mittermeier, C.G.; da Fonseca, G.A.B.; Kent, J. Biodiversity hotspots for conservation priorities. *Nature* **2000**, 403, 853–858. doi:10.1038/35002501. 448
11. Manning, J.; Goldblatt, P. *Plants of the Greater Cape Floristic Region. 1: The Core Cape flora*; South African National Biodiversity Institute, 2012. 449
12. Skowno, A.J.; Poole, C.J.; Raimondo, D.C. *National biodiversity assessment 2018: the status of South Africa's ecosystems and biodiversity. Synthesis Report*; South African National Biodiversity Institute, Pretoria, 2019. 450
13. Skowno, A.L.; Jewitt, D.; Slingsby, J.A. Rates and patterns of habitat loss across South Africa's vegetation biomes. *South African Journal of Science* **2021**, 117. doi:10.17159/sajs.2021/8182. 451
14. Humphreys, A.M.; Govaerts, R.; Ficinski, S.Z.; Lughadha, E.N.; Vorontsova, M.S. Global dataset shows geography and life form predict modern plant extinction and rediscovery. *Nature Ecology & Evolution* **2019**, 3, 1043–1047. doi:10.1038/s41559-019-0906-2. 452
15. Raimondo, D. The Red List of South African plants: a global first. *South African Journal of Science* **2011**, 107, 01–02. doi:10.4102/sajs.v107i3/4.653. 453
16. Brummitt, N.A.; Bachman, S.P.; Griffiths-Lee, J.; Lutz, M.; Moat, J.F.; Farjon, A.; Donaldson, J.S.; Hilton-Taylor, C.; Meagher, T.R.; Albuquerque, S.; et al. Green Plants in the Red: A Baseline Global Assessment for the IUCN Sampled Red List Index for Plants. *PLOS ONE* **2015**, 10, e0135152. doi:10.1371/journal.pone.0135152. 454
17. Von Hase, A.; Rouget, M.; Maze, K.; Helme, N. A fine-scale conservation plan for Cape lowlands renosterveld: technical report. *Report CCU* **2003**, 2, 104. 455
18. Moncrieff, G.R. Locating and Dating Land Cover Change Events in the Renosterveld, a Critically Endangered Shrubland Ecosystem. *Remote Sensing* **2021**, 13, 834. doi:10.3390/rs13050834. 456
19. Verbesselt, J.; Zeileis, A.; Herold, M. Near real-time disturbance detection using satellite image time series. *Remote Sensing of Environment* **2012**, 123, 98–108. doi:10.1016/j.rse.2012.02.022. 457
20. Zhu, Z.; Woodcock, C.E. Continuous change detection and classification of land cover using all available Landsat data. *Remote Sensing of Environment* **2014**, 144, 152–171. doi:10.1016/j.rse.2014.01.011. 458
21. Ye, S.; Rogan, J.; Zhu, Z.; Eastman, J.R. A near-real-time approach for monitoring forest disturbance using Landsat time series: stochastic continuous change detection. *Remote Sensing of Environment* **2021**, 252, 112167. doi:10.1016/j.rse.2020.112167. 459
22. Zhu, Z.; Zhang, J.; Yang, Z.; Aljaddani, A.H.; Cohen, W.B.; Qiu, S.; Zhou, C. Continuous monitoring of land disturbance based on Landsat time series. *Remote Sensing of Environment* **2020**, 238, 111116. doi:10.1016/j.rse.2019.03.009. 460
23. Kennedy, R.E.; Townsend, P.A.; Gross, J.E.; Cohen, W.B.; Bolstad, P.; Wang, Y.Q.; Adams, P. Remote sensing change detection tools for natural resource managers: Understanding concepts and tradeoffs in the design of landscape monitoring projects. *Remote sensing of environment* **2009**, 113, 1382–1396. 461
24. Hansen, M.C.; Potapov, P.V.; Moore, R.; Hancher, M.; Turubanova, S.A.; Tyukavina, A.; Thau, D.; Stehman, S.V.; Goetz, S.J.; Loveland, T.R. High-resolution global maps of 21st-century forest cover change. *science* **2013**, 342, 850–853. 462
25. Hansen, M.C.; Krylov, A.; Tyukavina, A.; Potapov, P.V.; Turubanova, S.; Zutta, B.; Ifo, S.; Margono, B.; Stolle, F.; Moore, R. Humid tropical forest disturbance alerts using Landsat data **2016**. 11, 034008. doi:10.1088/1748-9326/11/3/034008. 463
26. Kennedy, R.E.; Yang, Z.; Cohen, W.B. Detecting trends in forest disturbance and recovery using yearly Landsat time series: 1. LandTrendr — Temporal segmentation algorithms. *Remote Sensing of Environment* **2010**, 114, 2897–2910. doi:10.1016/j.rse.2010.07.008. 464

27. Pacheco-Pascagaza, A.M.; Gou, Y.; Louis, V.; Roberts, J.F.; Rodríguez-Veiga, P.; da Conceição Bispo, P.; Espírito-Santo, F.D.B.; Robb, C.; Upton, C.; Galindo, G.; et al. Near Real-Time Change Detection System Using Sentinel-2 and Machine Learning: A Test for Mexican and Colombian Forests. *Remote Sensing* **2022**, *14*, 707. doi:10.3390/rs14030707. 496
28. Wessels, K.J.; Van den Bergh, F.; Roy, D.P.; Salmon, B.P.; Steenkamp, K.C.; MacAlister, B.; Swanepoel, D.; Jewitt, D. Rapid land cover map updates using change detection and robust random forest classifiers. *Remote sensing* **2016**, *8*, 888. 497
29. Habib, T.; Inglada, J.; Mercier, G.; Chanussot, J. Support vector reduction in SVM algorithm for abrupt change detection in remote sensing. *IEEE Geoscience and Remote Sensing Letters* **2009**, *6*, 606–610. 498
30. Mas, J.F. Monitoring land-cover changes: a comparison of change detection techniques. *International journal of remote sensing* **1999**, *20*, 139–152. 499
31. Ismail Fawaz, H.; Forestier, G.; Weber, J.; Idoumghar, L.; Muller, P.A. Deep learning for time series classification: a review. *Data mining and knowledge discovery* **2019**, *33*, 917–963. 500
32. Kussul, N.; Lavreniuk, M.; Skakun, S.; Shelestov, A. Deep learning classification of land cover and crop types using remote sensing data. *IEEE Geoscience and Remote Sensing Letters* **2017**, *14*, 778–782. 501
33. Rußwurm, M.; Körner, M. Self-attention for raw optical satellite time series classification. *arXiv preprint arXiv:1910.10536* **2019**. 502
34. Pelletier, C.; Webb, G.I.; Petitjean, F. Temporal Convolutional Neural Network for the Classification of Satellite Image Time Series. *Remote Sensing* **2019**, *11*, 523. doi:10.3390/rs11050523. 503
35. Seifin, O.; Riese, F.M.; Keller, S. Deep learning for land cover change detection. *Remote Sensing* **2020**, *13*, 78. 504
36. Mucina, L.; Rutherford, M.C., Eds. *The vegetation of South Africa, Lesotho and Swaziland.*; Strelitzia 19. South African National Biodiversity Institute, Pretoria., 2006. 505
37. Rußwurm, M.; Pelletier, C.; Zollner, M.; Lefèvre, S.; Körner, M. BreizhCrops: A Time Series Dataset for Crop Type Mapping. *arXiv:1905.11893 [cs, stat]* **2020**. arXiv: 1905.11893. 506
38. Simoes, R.; Camara, G.; Queiroz, G.; Souza, F.; Andrade, P.R.; Santos, L.; Carvalho, A.; Ferreira, K. Satellite Image Time Series Analysis for Big Earth Observation Data. *Remote Sensing* **2021**, *13*, 2428. doi:10.3390/rs13132428. 507
39. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, \.; Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 2017, pp. 5998–6008. 508
40. Bahdanau, D.; Cho, K.; Bengio, Y. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv:1409.0473 [cs, stat]* **2016**. arXiv: 1409.0473. 509
41. Gorelick, N.; Hancher, M.; Dixon, M.; Ilyushchenko, S.; Thau, D.; Moore, R. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote sensing of Environment* **2017**, *202*, 18–27. 510
42. Chattopadhyay, A.; Sarkar, A.; Howlader, P.; Balasubramanian, V.N. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. 2018 IEEE winter conference on applications of computer vision (WACV). IEEE, 2018, pp. 839–847. 511
43. Tang, X.; Bullock, E.L.; Olofsson, P.; Estel, S.; Woodcock, C.E. Near real-time monitoring of tropical forest disturbance: New algorithms and assessment framework. *Remote Sensing of Environment* **2019**, *224*, 202–218. 512
44. Bullock, E.L.; Woodcock, C.E.; Holden, C.E. Improved change monitoring using an ensemble of time series algorithms. *Remote Sensing of Environment* **2020**, *238*, 111165. doi:10.1016/j.rse.2019.04.018. 513
45. Zhou, Q.; Rover, J.; Brown, J.; Worstell, B.; Howard, D.; Wu, Z.; Gallant, A.L.; Rundquist, B.; Burke, M. Monitoring Landscape Dynamics in Central U.S. Grasslands with Harmonized Landsat-8 and Sentinel-2 Time Series Data. *Remote Sensing* **2019**, *11*, 328. doi:10.3390/rs11030328. 514
46. Ye, S.; Rogan, J.; Zhu, Z.; Hawbaker, T.J.; Hart, S.J.; Andrus, R.A.; Meddens, A.J.H.; Hicke, J.A.; Eastman, J.R.; Kulakowski, D. Detecting subtle change from dense Landsat time series: Case studies of mountain pine beetle and spruce beetle disturbance. *Remote Sensing of Environment* **2021**, *263*, 112560. doi:10.1016/j.rse.2021.112560. 515
47. Parr, C.L.; Lehmann, C.E.; Bond, W.J.; Hoffmann, W.A.; Andersen, A.N. Tropical grassy biomes: misunderstood, neglected, and under threat. *Trends in ecology & evolution* **2014**, *29*, 205–213. 516
48. Rußwurm, M.; Wang, S.; Korner, M.; Lobell, D. Meta-learning for few-shot land cover classification. *Proceedings of the ieee/cvf conference on computer vision and pattern recognition workshops*, 2020, pp. 200–201. 517
49. Sun, X.; Wang, B.; Wang, Z.; Li, H.; Li, H.; Fu, K. Research progress on few-shot learning for remote sensing image interpretation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **2021**, *14*, 2387–2402. 518
50. Vanschoren, J. Meta-learning. In *Automated Machine Learning*; Springer, Cham, 2019; pp. 35–61. 519
51. Finn, C.; Abbeel, P.; Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. *International Conference on Machine Learning*. PMLR, 2017, pp. 1126–1135. 520
52. Maus, V.; Câmara, G.; Cartaxo, R.; Sanchez, A.; Ramos, F.M.; De Queiroz, G.R. A time-weighted dynamic time warping method for land-use and land-cover mapping. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **2016**, *9*, 3729–3739. 521
53. Dempster, A.; Petitjean, F.; Webb, G.I. ROCKET: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining and Knowledge Discovery* **2020**, *34*, 1454–1495. 522