Supervised continuous land cover change detection in a critically endangered shrubland ecosystem

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

Existing efforts to rapidly detect land cover change in 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 allow for unsupervised methods, often with manually determined thresholds, to detect land cover change 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 supervised 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 training dataset of precisely dated vegetation change events between 2016 and 2020 was obtained from daily, high resolution Planet labs satellite data. This dataset was then used to train 1D convolutional neural networks and Transformers to continuously classify land cover change events in multivariate time-series of vegetation activity from Sentinel 2 satellites as new data becomes available. These models reached a f-score of 0.93, a 61% improvement over the f-score of 0.57 achieved using an unsupervised method designed for forested ecosystems. 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 labelled datasets of land cover change events with which to train machine learning classifiers.

Land cover change | Land cover monitoring | Deep learning | Renosterveld | Threatened ecosystems | Sentinel 2 | Planet labs

Correspondence: ***glenn@saeon.ac.za***

Introduction

Land cover change is the single largest cause underlying the global biodiversity crisis (Brondizio et al., 2019). It is estimated that at least 21% of all terrestrial plant extinctions are directly related to agriculture alone (Le Roux et al., 2019). 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 (Turner, 2014; Turner et al., 2003; Yang et al., 2013). 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 (Bond, 2019; Poulter et al., 2014)). 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 (Slingsby et al., 2020; Wilson et al., 2015). 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 (Manning and Goldblatt, 2012; Myers et al., 2000). Within the CFR the most threatened vegetation type is the grassy shrubland known as Renosterveld (Skowno et al., 2019, 2021). Renosterveld is restricted to relativity 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 (Brummitt et al., 2015; Humphreys et al., 2019; Raimondo, 2011; Skowno et al., 2021). Only 5-10% its original extent remains, with much of this contained in small isolated fragments (Von Hase et al., 2003). Despite legislation prohibiting the destruction

of Renosterveld habitat, loss is ongoing (Moncrieff, 2021). It is estimated that at current rates Renosterveld will become extinct as an ecosystem within the 21st century (Skowno et al., 2021). 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.

The current approach to conservation in the region is clearly inadequate. 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 - thereby deterring future habitat destruction. The identification of habitat loss through the direct comparison of images taken taken months or years apart can produce very accurate maps of change within a discrete time period. A range of methods are available to achieve this ranging from the very simple (e.g. direct difference of vegetation indexes) to supervised methods that rely on validated ground truth data combined with machine learning approaches (Kennedy et al., 2009). Moncrieff (2021) used a training dataset of validated land cover change events to produce a map of Renosterveld loss between 2016 and 2020 after manually verifying potential change identified through model predictions.

While the direct comparison of paired images provides useful information for monitoring and reporting on habitat loss, this information is usually only obtained months or years after the reported change has occurred (Kennedy et al., 2010). 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 building a model that can accurately predict the vegetation activity of intact forests, and comparing predictions to recently acquired observations (Verbesselt et al., 2012; Ye et al., 2021a; Zhu and Woodcock, 2014; Zhu et al., 2020). When observed vegetation activity falls outside out the range expected for natural vegetation, this is taken to be indicative of forest loss. A variety of methods have been used to model expected vegetation activity, ranging from neural networks to relativity simple linear models (Tang et al., 2019). The thresholds used to determine whether anomalies are indeed true change are often predetermined heuristically. This is necessary as determining the degree of change in vegetation activity that is indicative of vegetation cover change from data requires not only maps of where change has occurred, but also precise dates of the change events. 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. This problem is particularly acute in non-forest ecosystems with natural dynamics such as fire, aridity and herbivory (Slingsby et al., 2020).

If ground-truth data are available, time-series of vegetation activity can be classified into land cover types or land cover change using supervised learning. This approach is still limited to producing updates on land cover change at fixed intervals (typically annually), as the training data typically do not specify the exact date on which change occurred, but rather that change occurred between two fixed points in time (e.g. sometime between January and December). Thus a new full year of data are required before a new prediction can be made. If a precise date for a land cover change event is available, time series classifiers that are robust to seasonal variation and natural variability can be trained (Figure 1). 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 unsupervised change detection methods, as the models could theoretically learn to distinguish the spectral patterns indicative of land cover change from natural disturbances. Hansen et al. (2016) 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. Approaches using neural network architectures such as convolutional neural networks (CNN) or self-attention have shown improved performance on land cover classification tasks relative to more traditional tree-based machine learning approaches, but are not yet widely applied to continuous change detection (Kussul et al., 2017; Pelletier et al., 2019; Rußwurm and Körner, 2019).

Here I describe the design and performance of continuous land cover change detection in a critically endangered shrubland ecosystem, the Renosterveld of South Africa, based on supervised machine learning approaches. Using an dataset of precisely dated land cover change events, a variety of deep learning architectures are trained and evaluated against tree-based and unsupervised methods commonly used for land cover change detection and monitoring in forested ecosystems. Models building and design choices are motivated by an intention to produce a system that can be deployed for operationally use by conservation management and law enforcement agencies.

Methods

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 (Mucina and Rutherford, 2006). 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 of 12,241 km^2, with Renosterveld covering 4711 km^2 with this region naturally. Detailed mapping of remaining natural vegetation in the Overberg based on 2003 data identified 656 km^2 of the original extent of Renosterveld vegetation in the Overberg remaining.

Using supervised change detection on Sentinel 2 imagery and

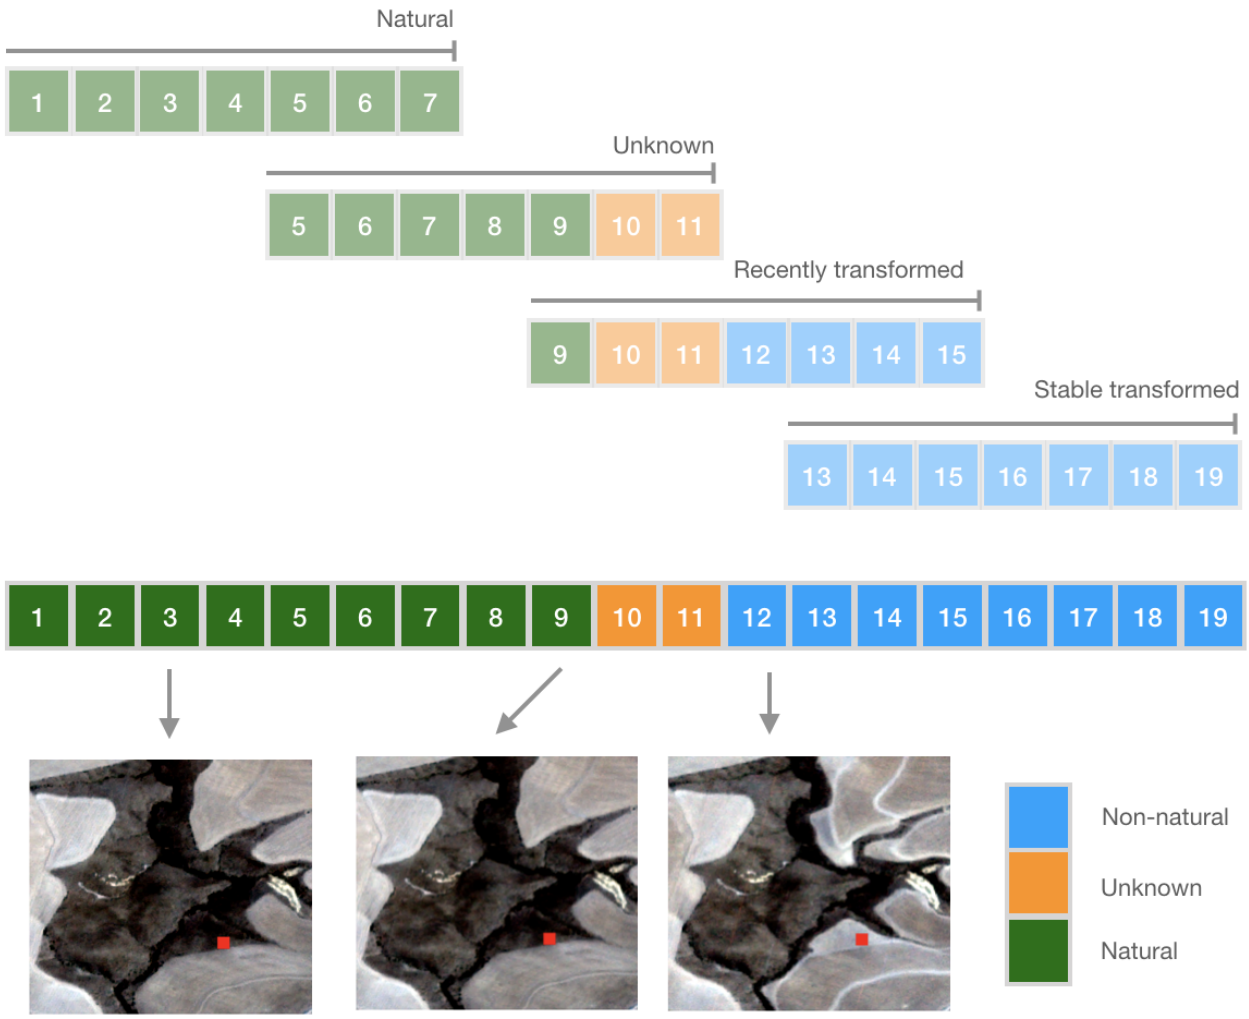

Fig. 1. The approach used for assigning labels to time series. Each step in the sequence of 19 observations represents an observation obtained for a single pixel at a regular interval (e.g. 10 days). The land cover label for the pixel (identified by the red point in the images) can either be natural, 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.

subsequent detailed verification with very high resolution imagery, Moncrieff (2021) identified all conversion of Renosterveld to non-natural land cover between January 2016 and January 2020. Land cover change was identified within 268 spatially contiguous parcels 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. These two dates occur within one week for 21 percent of all events, and within two months for 87 percent. 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. 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 the no change data parcels of vegetation that remained natural over the entire study period were added. This added an additional 2590 ha to the training and validation dataset and 24 ha to the testing dataset. These no-change parcels were selected from the (Von Hase et al., 2003) map of remaining Renosterveld, buffered to ensure that they did not occur within 100m of non-natural vegetation and manually verified.

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 preprocessed. 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: NDVI, EVI, NDWI, NBR, 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 become available at least every 5 days subsequent to the launch of Sentinel 2B. In parts of the study region data is available at frequencies higher than 5 days due to swath overlap. To create regular time series from the irregularly spaced 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, the reflectance and vegetation indices 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'.

The time series for each pixel 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 1). 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 cove 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 occur before the start of the time series, land cover change has already taken place. Here there is no need to identify stable non-natural land cover, thus these data are discarded.

The resulting labelled dataset is highly imbalanced, with the vast majority of samples in the 'natural' class. There is also very high spatial auto-correlation among neighbouring pixels within the same parcel. Successive time series for the same pixel obtained 10 days apart will share 94 percent of observations. To ameliorate class imbalance data 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. Subsequent to this data thinning, parcels are assigned to training and validation sets using a 70/30 split. To reduce temporal oversampling within a pixel, time series windows of 180 days are shifted forward by 60 days for every sample taken. This represents a trade off between reducing temporal oversampling and discarding labelled data. Test data does not undergo temporal or spatial thinning, and a new prediction is made every 10 days to quantify performance in an operational setting.

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 (Pelletier et al., 2019). 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 (Rußwurm et al., 2020a; Simoes et al., 2021).

Transformer models were originally proposed as sequence-to-sequence encoder-decoder models based on the self-attention operation and applied to language translation (Vaswani et al., 2017). Attention mechanisms operate by assigning an importance-score (attention) to weight each element in a sequence (Bahdanau et al., 2016). 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. In this work, a Transformer model that has been evaluated in (Rußwurm and Körner, 2019) 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 on satellite image time series (Rußwurm et al., 2020a).

The accurately dated land cover change data used here to train supervised methods is seldom available. Unsupervised approaches to land cover change detection are more generally applicable and widely used (Bullock et al., 2020; Tang et al., 2019). A simple unsupervised baseline is used to evaluate the improvement in accuracy achieved when using supervised methods on these data. The approach followed to construct the unsupervised baseline is very similar to the approached outlined by Zhu and Woodcock (2014), 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 unsupervised methods that would improve on this baseline are available (Ye et al., 2021a,b; Zhou et al., 2019), but they are not commonly used on S2 data.

The unsupervised baseline operates by fitting a simple linear model with a single harmonic term to capture seasonal variation using robust regression through iteratively reweighted

Table 1. Model performance evaluated on test data. Recall-1-short reports model recall when change has occured 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-short
Random Forest	0.98	0.88	0.79	**0.99**	**0.98**	0.98	
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	0.93	0.57	0.48	0.98	0.70	0.94	-

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 exceed the prediction by 3 times the RMSE, land cover change is assumed to have occurred. More sophisticated implementations of this workflow make use of multiple bands, or use parameters learned from labelled land cover change events (Ye et al., 2021b).

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 is minimize focal loss. TempCNN and Transfomers 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 (Gorelick et al., 2017).

To aid model interpretation, the most important temporal regions identified by the TempCNN model for detecting land cover change are calculated using Grad-CAM++ (Chattopadhay et al., 2018). Originally developed for use with CNNs and image classification, Grad-CAM++ can easily be used when CNNs are applied to time series classification. The importance of each time-step in the classification of a time series is calculated using a weighted average of all positive activation gradients across channels in the final convolutional layer.

Results

The best performing 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 filter 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 best performing supervised model was the TempCNN, with a recall of 0.89 and precision of 0.96 when identifying Renosterveld loss 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 Renosterveld loss class and overall f-score of 0.92, while the Random Forest was significantly worse than both TempCNN and the Transfomer, with Renosterveld loss recall and precision of 0.79 and 0.98, and overall f-score of 0.88. All supervised models outperformed unsupervised classification, which was far more likely to produce both false negatives and false positives with recall of 0.48 and precision of 0.70 on the Renosterveld loss class and overall f-score of 0.57.

Performance for all supervised models declines when their ability to detect Renosterveld loss 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. Unsupervised 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 2 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 saliency of these temporal regions.

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/global-renosterveld-watch and an example sequence of predictions over 2 months in an area throughout course of a land cover change event is depicted in Figure 3. 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

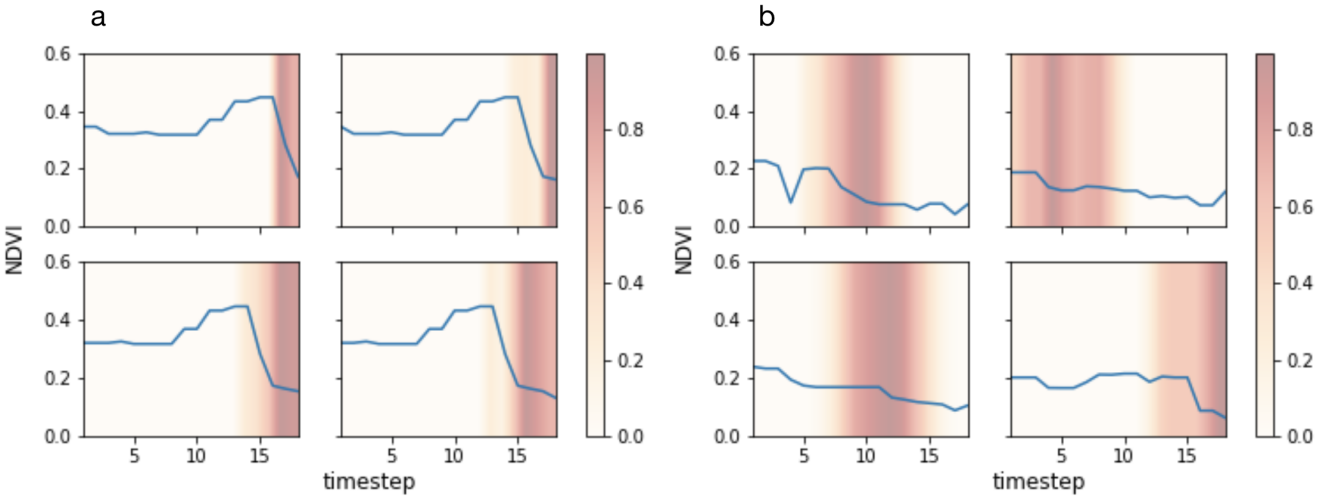

Fig. 2. Saliency of temporal regions in time series for pixels in which land cover change is predicted to have occurred, determined using Grad-CAM++. Higher saliency scores indicates that the region is more influential in determining that land cover change has occurred. Blue lines show NDVI, though saliency 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.

updating of existing regional land cover maps.

Discussion

Shrublands are globally important repositories of biodiversity and play an important role in regulating planetary biochemical cycles (Bond, 2019; Poulter et al., 2014). In many regions shrublands are being lost at a rate that exceeds other ecosystems such as forests (Parr et al., 2014). 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 unsupervised methods typically applied in operational forest monitoring perform poorly. This study demonstrates that reliable detection of land cover change in shrublands in 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 use of supervised classification, enabled by the availability of hand-labelled change events with the date of each event precisely determined.

Supervised models for continuous detection of land cover change are rare. The labels needed to train them 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 (Rußwurm et al., 2020b; Sun et al., 2021). 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 requirement further painstaking data collection and labelling - meta-learning approaches ought to reduce data requirements for application to new disturbance types and new geographies (Vanschoren, 2019). 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 (Finn et al., 2017; Rußwurm et al., 2020b).

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

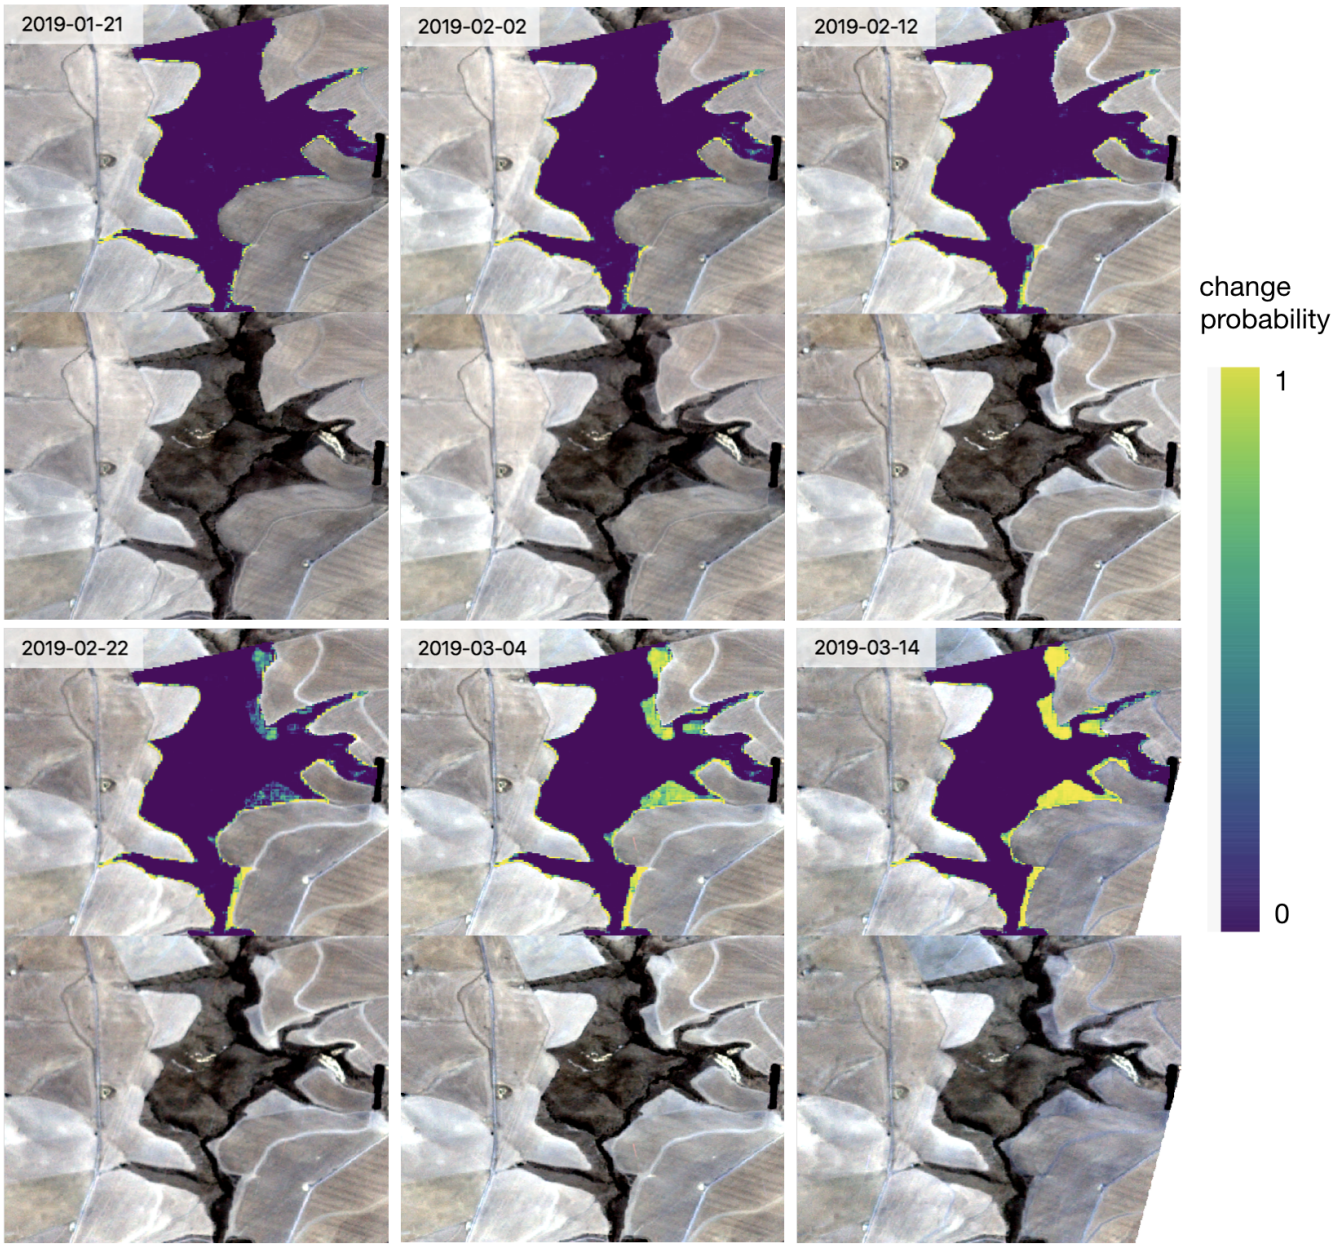

Fig. 3. Example sequence of predictions through the course of a land cover change event in Renosterveld shrublands of the Overberg region, South Africa. Upper panels show the predicted probability of land cover change, where predictions are only made in pixels known to have intact natural vegetation 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.

many intended applications of near-realtime continuous land cover change monitoring - such as law enforcement and management intervention. Unsupervised approaches typically require multiple consecutive observations confirming the change has occurred before a positive identification is made (Ye et al., 2021a; Zhu and Woodcock, 2014). 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 observation 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 prepossessing and no atmospheric correction is performed (Rußwurm and Körner, 2019; Rußwurm et al., 2020a). Here, extensive feature engineering has been conducted and cloud masking performed to remove noisy data, potentially result-

ing 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 (Dempster et al., 2020; Maus et al., 2016). 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 using tools such as Google Earth Engine and Google Cloud Dataflow.

Conclusions

The potential of supervised 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 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.

ACKNOWLEDGEMENTS

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. This research was supported 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, and the Group on Earth Observations-Google Earth Engine Programme. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

DATA AVAILABILITY

Code and data used in this analysis are available at https://github.com/GMoncrieff/renosterveld-monitor. Due to ongoing criminal investigations location data has been removed to protect the identity of landowners where necessary.

Bibliography

D. Bahdanau, K. Cho, and Y. Bengio. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv:1409.0473 [cs, stat]*, May 2016. URL http://arxiv.org/abs/1409.0473. arXiv: 1409.0473.

W. J. Bond. *Open ecosystems: ecology and evolution beyond the forest edge*. Oxford University Press, 2019.

E. S. Brondizio, J. Settele, S. Díaz, and H. T. Ngo. Global assessment report on biodiversity and ecosystem services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. 2019.

N. A. Brummitt, S. P. Bachman, J. Griffiths-Lee, M. Lutz, J. F. Moat, A. Farjon, J. S. Donaldson, C. Hilton-Taylor, T. R. Meagher, S. Albuquerque, E. Aletrari, A. K. Andrews, G. Atchison, E. Baloch, B. Barlozzini, A. Brunazzi, J. Carretero, M. Celesti, H. Chadburn, E. Cianfoni, C. Cockel, V. Coldwell, B. Concetti, S. Contu, V. Crook, P. Dyson, L. Gardiner, N. Ghanim, H. Greene, A. Groom, R. Harker, D. Hopkins, S. Khela, P. Lakeman-Fraser, H. Lindon, H. Lockwood, C. Loftus, D. Lombrici, L. Lopez-Poveda, J. Lyon, P. Malcolm-Tompkins, K. McGregor, L. Moreno, L. Murray, K. Nazar, E. Power, M. Q. Tuijtelaars, R. Salter, R. Segrott, H. Thacker, L. J. Thomas, S. Tingvoll, G. Watkinson, K. Wojtaszekova, and E. M. N. Lughadha. Green Plants in the Red: A Baseline Global Assessment for the IUCN Sampled Red List Index for Plants. *PLOS ONE*, 10(8):e0135152, Aug. 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0135152. URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0135152.

E. L. Bullock, C. E. Woodcock, and C. E. Holden. Improved change monitoring using an ensemble of time series algorithms. *Remote Sensing of Environment*, 238:111165, Mar. 2020. ISSN 0034-4257. doi: 10.1016/j.rse.2019.04.018. URL https://www.sciencedirect.com/science/article/pii/S0034425719301634.

A. Chattopadhay, A. Sarkar, P. Howlader, and V. N. Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE winter conference on applications of computer vision (WACV)*, pages 839–847. IEEE, 2018.

A. Dempster, F. Petitjean, and G. I. Webb. ROCKET: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining and Knowledge Discovery*, 34 (5):1454–1495, 2020.

C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.

N. Gorelick, M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote sensing of Environment*, 202:18–27, 2017.

M. C. Hansen, A. Krylov, A. Tyukavina, P. V. Potapov, S. Turubanova, B. Zutta, S. Ifo, B. Margono, F. Stolle, and R. Moore. Humid tropical forest disturbance alerts using Landsat data. 11(3): 034008, Mar. 2016. ISSN 1748-9326. doi: 10.1088/1748-9326/11/3/034008. URL https://doi.org/10.1088/1748-9326/11/3/034008.

A. M. Humphreys, R. Govaerts, S. Z. Ficinski, E. N. Lughadha, and M. S. Vorontsova. Global dataset shows geography and life form predict modern plant extinction and rediscovery. *Nature Ecology & Evolution*, 3(7):1043–1047, July 2019. ISSN 2397-334X. doi: 10.1038/s41559-019-0906-2. URL https://www.nature.com/articles/s41559-019-0906-2.

R. E. Kennedy, P. A. Townsend, J. E. Gross, W. B. Cohen, P. Bolstad, Y. Q. Wang, and P. Adams. Remote sensing change detection tools for natural resource managers: Understanding concepts and tradeoffs in the design of landscape monitoring projects. *Remote sensing of environment*, 113(7):1382–1396, 2009.

R. E. Kennedy, Z. Yang, and W. B. Cohen. Detecting trends in forest disturbance and recovery using yearly Landsat time series: 1. LandTrendr — Temporal segmentation algorithms. *Remote Sensing of Environment*, 114(12):2897–2910, Dec. 2010. ISSN 0034-4257. doi: 10.1016/j.rse.2010.07.008. URL http://www.sciencedirect.com/science/article/pii/S0034425710002245.

N. Kussul, M. Lavreniuk, S. Skakun, and A. Shelestov. Deep learning classification of land cover and crop types using remote sensing data. *IEEE Geoscience and Remote Sensing Letters*, 14(5):778–782, 2017.

J. J. Le Roux, C. Hui, M. L. Castillo, J. M. Iriondo, J.-H. Keet, A. A. Khapugin, F. Médail, M. Rejmánek, G. Theron, F. A. Yannelli, and H. Hirsch. Recent Anthropogenic Plant Extinctions Differ in Biodiversity Hotspots and Coldspots. *Current Biology*, 29(17):2912–2918.e2, Sept. 2019. ISSN 0960-9822. doi: 10.1016/j.cub.2019.07.063. URL https://www.sciencedirect.com/science/article/pii/S0960982219309431.

J. Manning and P. Goldblatt. *Plants of the Greater Cape Floristic Region. 1: The Core Cape flora*. South African National Biodiversity Institute, 2012.

V. Maus, G. Câmara, R. Cartaxo, A. Sanchez, F. M. Ramos, and G. R. De Queiroz. 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*, 9(8):3729–3739, 2016.

G. R. Moncrieff. Locating and Dating Land Cover Change Events in the Renosterveld, a Critically Endangered Shrubland Ecosystem. *Remote Sensing*, 13(5):834, Jan. 2021. doi: 10.3390/rs13050834. URL https://www.mdpi.com/2072-4292/13/5/834.

L. Mucina and M. C. Rutherford, editors. *The vegetation of South Africa, Lesotho and Swaziland*. Strelitzia 19. South African National Biodiversity Institute, Pretoria., 2006.

N. Myers, R. A. Mittermeier, C. G. Mittermeier, G. A. B. da Fonseca, and J. Kent. Biodiversity hotspots for conservation priorities. *Nature*, 403(6772):853–858, Feb. 2000. ISSN 1476-4687. doi: 10.1038/35002501. URL https://www.nature.com/articles/35002501.

C. L. Parr, C. E. Lehmann, W. J. Bond, W. A. Hoffmann, and A. N. Andersen. Tropical grassy biomes: misunderstood, neglected, and under threat. *Trends in ecology & evolution*, 29(4): 205–213, 2014.

C. Pelletier, G. I. Webb, and F. Petitjean. Temporal Convolutional Neural Network for the Classification of Satellite Image Time Series. *Remote Sensing*, 11(5):523, Jan. 2019. doi: 10.3390/rs11050523. URL https://www.mdpi.com/2072-4292/11/5/523.

B. Poulter, D. Frank, P. Ciais, R. B. Myneni, N. Andela, J. Bi, G. Broquet, J. G. Canadell, F. Chevallier, Y. Y. Liu, S. W. Running, S. Sitch, and G. R. van der Werf. Contribution of semi-arid ecosystems to interannual variability of the global carbon cycle. *Nature*, 509 (7502):600–603, May 2014. ISSN 1476-4687. doi: 10.1038/nature13376. URL https://www.nature.com/articles/nature13376.

D. Raimondo. The Red List of South African plants: a global first. *South African Journal of Science*, 107(3-4):01–02, Apr. 2011. ISSN 0038-2353. doi: 10.4102/sajs.v107i3/4.653. URL http://www.scielo.org.za/scielo.php?script=sci_abstract&pid=S0038-23532011000200003&lng=en&nrm=iso&tlng=en.

M. Rußwurm and M. Körner. Self-attention for raw optical satellite time series classification. *arXiv preprint arXiv:1910.10536*, 2019.

M. Rußwurm, C. Pelletier, M. Zollner, S. Lefèvre, and M. Körner. BreizhCrops: A Time Series Dataset for Crop Type Mapping. *arXiv:1905.11893 [cs, stat]*, May 2020a. URL http://arxiv.org/abs/1905.11893. arXiv: 1905.11893.

M. Rußwurm, S. Wang, M. Korner, and D. Lobell. Meta-learning for few-shot land cover classifi-

cation. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition workshops*, pages 200–201, 2020b.

R. Simoes, G. Camara, G. Queiroz, F. Souza, P. R. Andrade, L. Santos, A. Carvalho, and K. Ferreira. Satellite Image Time Series Analysis for Big Earth Observation Data. *Remote Sensing*, 13(13):2428, Jan. 2021. doi: 10.3390/rs13132428. URL https://www.mdpi.com/2072-4292/13/13/2428.

A. J. Skowno, C. J. Poole, and D. C. Raimondo. *National biodiversity assessment 2018: the status of South Africa's ecosystems and biodiversity. Synthesis Report*. South African National Biodiversity Institute, Pretoria, 2019.

A. L. Skowno, D. Jewitt, and J. A. Slingsby. Rates and patterns of habitat loss across South Africa's vegetation biomes. *South African Journal of Science*, 117(1/2), Jan. 2021. ISSN 1996-7489. doi: 10.17159/sajs.2021/8182. URL https://sajs.co.za/article/view/8182.

J. A. Slingsby, G. R. Moncrieff, and A. M. Wilson. Near-real time forecasting and change detection for an open ecosystem with complex natural dynamics. *ISPRS Journal of Photogrammetry and Remote Sensing*, 166:15–25, Aug. 2020. ISSN 0924-2716. doi: 10.1016/j.isprsjprs.2020.05.017. URL http://www.sciencedirect.com/science/article/pii/S092427162030143X.

X. Sun, B. Wang, Z. Wang, H. Li, H. Li, and K. Fu. Research progress on few-shot learning for remote sensing image interpretation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:2387–2402, 2021.

X. Tang, E. L. Bullock, P. Olofsson, S. Estel, and C. E. Woodcock. Near real-time monitoring of tropical forest disturbance: New algorithms and assessment framework. *Remote Sensing of Environment*, 224:202–218, 2019.

W. Turner. Sensing biodiversity. *Science*, 346(6207):301–302, 2014.

W. Turner, S. Spector, N. Gardiner, M. Fladeland, E. Sterling, and M. Steininger. Remote sensing for biodiversity science and conservation. *Trends in ecology & evolution*, 18(6):306–314, 2003.

J. Vanschoren. Meta-learning. In *Automated Machine Learning*, pages 35–61. Springer, Cham, 2019.

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, \. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

J. Verbesselt, A. Zeileis, and M. Herold. Near real-time disturbance detection using satellite image time series. *Remote Sensing of Environment*, 123:98–108, Aug. 2012. ISSN 0034-4257. doi: 10.1016/j.rse.2012.02.022. URL https://www.sciencedirect.com/science/article/pii/S0034425712001150.

A. Von Hase, M. Rouget, K. Maze, and N. Helme. A fine-scale conservation plan for Cape lowlands renosterveld: technical report. *Report CCU*, 2(03):104, 2003.

A. M. Wilson, A. M. Latimer, and J. A. J. Silander. 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*, 112(29):9058–9063, 2015. ISSN 1091-6490 (Electronic). doi: 10.1073/pnas.1416710112.

J. Yang, P. Gong, R. Fu, M. Zhang, J. Chen, S. Liang, B. Xu, J. Shi, and R. Dickinson. The role of satellite remote sensing in climate change studies. *Nature climate change*, 3(10):875–883, 2013.

S. Ye, J. Rogan, Z. Zhu, and J. R. Eastman. A near-real-time approach for monitoring forest disturbance using Landsat time series: stochastic continuous change detection. *Remote Sensing of Environment*, 252:112167, Jan. 2021a. ISSN 0034-4257. doi: 10.1016/j.rse.2020.112167. URL https://www.sciencedirect.com/science/article/pii/S003442572030540X.

S. Ye, J. Rogan, Z. Zhu, T. J. Hawbaker, S. J. Hart, R. A. Andrus, A. J. H. Meddens, J. A. Hicke, J. R. Eastman, and D. Kulakowski. Detecting subtle change from dense Landsat time series: Case studies of mountain pine beetle and spruce beetle disturbance. *Remote Sensing of Environment*, 263:112560, Sept. 2021b. ISSN 0034-4257. doi: 10.1016/j.rse.2021.112560. URL https://www.sciencedirect.com/science/article/pii/S0034425721002807.

Q. Zhou, J. Rover, J. Brown, B. Worstell, D. Howard, Z. Wu, A. L. Gallant, B. Rundquist, and M. Burke. Monitoring Landscape Dynamics in Central U.S. Grasslands with Harmonized Landsat-8 and Sentinel-2 Time Series Data. *Remote Sensing*, 11(3):328, Jan. 2019. doi: 10.3390/rs11030328. URL https://www.mdpi.com/2072-4292/11/3/328.

Z. Zhu and C. E. Woodcock. Continuous change detection and classification of land cover using all available Landsat data. *Remote Sensing of Environment*, 144:152–171, Mar. 2014. ISSN 0034-4257. doi: 10.1016/j.rse.2014.01.011. URL https://www.sciencedirect.com/science/article/pii/S0034425714000248.

Z. Zhu, J. Zhang, Z. Yang, A. H. Aljaddani, W. B. Cohen, S. Qiu, and C. Zhou. Continuous monitoring of land disturbance based on Landsat time series. *Remote Sensing of Environment*, 238:111116, Mar. 2020. ISSN 0034-4257. doi: 10.1016/j.rse.2019.03.009. URL https://www.sciencedirect.com/science/article/pii/S0034425719301002.