Operational Google-Earth-Engine Workflow to Monitor Irrigated Areas in a Semi-Arid Climate

Kilian Vos, Sikdar M. M. Rasel, Azeem Khan, Yi Lu, Mustak Shaikh, Vaibhav Gupta, Michelle Newstead

Water Group, Department of Climate Change, Energy, the Environment and Water (DCCEEW), Government of New South Wales, Australia − kilian.vos, sikdar.rasel, azeem.khan, yi.lu, mustak.shaikh, vaibhav.gupta, michelle.newstead@dpie.nsw.gov.au

Keywords: Sentinel-2, Landsat, Water Resources, Irrigated Areas, Automation, Cloud Computing.

Abstract

Sustainable water resources management relies on advanced hydrological models calibrated from long-term satellite observations of water cycles. Monitoring water resources from Earth-observing satellites is enabled by cloud computing environments like Google Earth Engine (GEE) that support easy processing and analysis of archived data. While there are different satellite-based indicators of water resources availability, this paper focuses on mapping irrigated areas for each summer harvest season. The paper presents an automated workflow to extract yearly irrigated areas from imagery made public by Landsat and Sentinel-2. This method uses a normalised difference vegetation index (NDVI) to identify fields that exhibit higher vegetation greenness than their surroundings and are potentially being irrigated. As a first step, Google Earth Engine was used to create a seasonal maximum NDVI composite on which a grid-based thresholding algorithm was applied to identify irrigated areas in different parts of a catchment. The irrigated areas were then refined using morphological image processing and a quality-control step was conducted with ancillary datasets such as the most recent landuse classification, evapotranspiration data and crop phenology information. Finally, real-world use cases are presented to demonstrate the ability of satellite remote sensing to provide critical information on irrigation practices over the last 4 decades at the catchment scale.

1. Introduction

1.1. Background

Climate change and global warming are exacerbating water scarcity, resulting in increased aridity and shifts in agroecological zones that affect crop yields. Although only 18% of cultivated areas globally are irrigated, 70% of available freshwater resources are used for irrigation (Magidi et al., 2015). These resources play a significant role in semi-arid climates such as those experienced in much of Australia. The Australian continent is the sixth largest country in the world and covers a large range of climatic zones, from the tropics in the north to the arid interior and temperate regions in the south. In terms of rainfall, Australia is the driest inhabited continent. The amount of rainwater that flows into rivers is also relatively low – on average only 12% of rainfall flows into rivers, compared to 39% in Europe and 52% in North America. Additionally, rainfall patterns at the inter-annual timescale are controlled by the state of the El Niño-Southern Oscillation (ENSO) climate phenomenon. Long-term trends of precipitation show a shift towards drier conditions in the southwest and southeast regions, with a higher frequency of years of below average rainfall in recent decades. In this context, freshwater resources are critical to support industry (agriculture and mining), communities (drinking water) and environmental assets such as wetlands. Periods of water scarcity present particular challenges for agriculture, which accounts for approximately 75% of Australia’s fresh water use. In response, farmers are implementing innovative technologies like precision farming, water-efficient technologies, and enhanced water storage facilities to optimise water allocation and minimise wastage.

However, to ensure better water resources outcomes for NSW, continuous monitoring of irrigated water usage is also important. In theory water planners and managers could rely on *in situ* metering devices to do this. However, given the extent of irrigated lands in NSW, this is not economically or practically feasible for large areas over long time periods. Satellite remote sensing offers a low-cost solution, leveraging publicly available imagery and recent advances in cloud computing platforms to monitor water usage across large spatio-temporal scales. In particular, the extent of irrigated areas each season is observable from satellite and can be used to estimate how much water is being used in agricultural areas. This information could then inform targeted policy interventions to improve the long-term sustainability of water resources for the benefit of all stakeholders.

Multiple irrigated areas products have been developed in the past, including the Global Map of Irrigated Areas version 5 (GMIA 5.0), the MIRCA 2000 product, and the International Water Management Institute (IWMI) Irrigated Area Map. These have used varying spatial scales and resolutions. Various spectral indices and algorithms are also used for satellite-based irrigated-area mapping. Among those indices, the Normalised Difference Vegetation Index (NDVI) is endorsed as the most suitable measure of canopy greenness, based on its sensitivity to plant canopy chlorophyll content (Rouse et al., 1973). Near-infrared (NIR) (0.85-0.89 µm) radiation used in NDVI is scattered by the physical structure of the leaf, resulting in high levels of reflectance and transmittance, while NDVI’s red (0.63-0.69 µm) radiation is strongly absorbed by leaf pigments involved in photosynthesis (Jensen, 2021). Cultivated crops exhibit significant changes in pigment concentration and structure as they develop from the early growing season through to maturity and senescence. Such changes affect reflectance in NIR and red radiation, thus providing the theoretical basis for NDVI to use reflected radiation from satellite images to retrieve a crop-specific growth patterns/optical signature (including sowing,

tillering, heading and maturity points), along with crop age and any crop stress.

NDVI values range from -1 to +1, where higher values correspond to greener and denser vegetation. Conversely, negative NDVI values typically represent non-vegetated surfaces such as water bodies and urban areas. In the life cycle of plants, the NDVI profile typically rises during the vegetative growth stage, reaching a peak followed by a plateau once the plant reaches maturity. Subsequently, as plants undergo senescence, the NDVI profile gradually declines. Therefore, the timing of vegetation ‘green-up ’and senescence can be determined using time series of NDVI data. Several methods exist that use NDVI thresholds, backward-looking moving-averages or empirical equations.

When advances in remote sensing data (multispectral and hyperspectral), processing methods and techniques are combined with those in big data and cloud computing, irrigated areas can be more easily classified. However, although hyperspectral imagery provides more phenological information, there are challenges associated with data capture and processing. First, tasking an aerial capture with a hyperspectral sensor can be prohibitively expensive. Second, it is subject to seasonal, illumination and aviation restrictions. In addition, processing and extracting information from high multi-dimensional data containing hundreds of narrow continuous bands (Rasel et al., 2019) is very challenging.

Benefits of multispectral satellite imagery are that it is publicly available and provides regular observations at medium spatial resolution (10 m for Sentinel-2 and 30 m for Landsat). Moreover, advances in classification algorithms, including Random Forest, Support Vector Machine, Decision Trees and Extreme Gradient Boosting (Abdi, 2020), make it possible to efficiently extract features from multispectral satellite imagery. Cloud computing platforms like Google Earth Engine (GEE) also provide access to petabytes of remote-sensing data and enable users to deploy remote sensing algorithms on the cloud for large-scale processing (Teleguntla et al,. 2018).

The aim of this article is to present an operational GEE-enabled workflow to generate long-term time-series of irrigated area maps from Landsat and Sentinel-2. Section 2 presents the study site and methodology, while Section 3 discusses real-world applications of the satellite-derived irrigated areas within the water resources management sector.

1.2. Novelty and Context of this Article

Although previous research has investigated the creation of regional-scale irrigated area maps from Earth-observing satellites (Magidi et al., 2015, Wardlow et al., 2008), there is limited research demonstrating an operational workflow for irrigated areas mapping deployed in the cloud and used for real-world applications. This cloud-based workflow provides efficient retrieval of satellite imagery and ancillary datasets (evapotranspiration, rainfall, landuse) from multiple sources inside one platform. In addition, this paper proposes a grid-based thresholding algorithm that accounts for spatial variability in the landscape across an entire catchment. For example, irrigated areas and native vegetation will have different spectral signatures in different parts of a catchment based on elevation, climate and other geographic factors. In this workflow, each component can be fine-tuned based on the user requirements and adapted to different geographical areas. In summary, the approach outlined in this paper:

1. automatically generates irrigated crop area maps from GEE based on Landsat and Sentinel-2 imagery;
2. refines the irrigated area maps using ancillary datasets such as evapotranspiration data and recent landuse classification;
3. provides a framework to validate irrigated areas that is up to industry standard and can be used in subsequent applications in water resources management.

2. Data and Methodology

2.1. Study Site

The study area (Figure 1) is the agricultural catchment of Macquarie-Castlereagh in the Murray Darling Basin (MDB), also known as the ‘Food Bowl of Australia’.

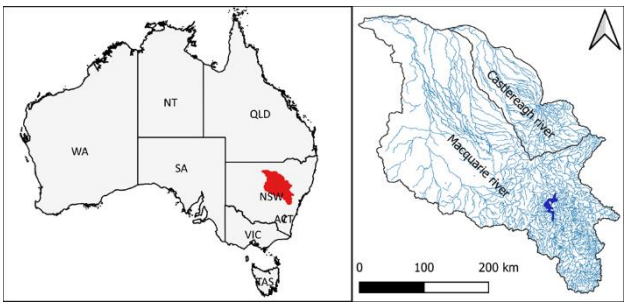

Figure 1. Location of Macquarie and Castlereagh Catchment in New South Wales, Australia.

2.2. Satellite Imagery and Evapotranspiration Data

2.2.1. Satellite imagery

Medium-resolution and cloud-free satellite images from Sentinel 2 (10 m/pixel) from 2015 to 2023 and Landsat 5, 7 and 8 imagery (30 m/pixel) from 1985 to 2014 were analysed in Google Earth Engine (GEE) cloud computing environment.

2.2.2. Evapotranspiration (ET)

This is raster data that measures the loss of water from both plants and soil and provides an indication for irrigation. The CSIRO CMRSET provides monthly-averaged actual evapotranspiration data (in mm/day) for Australia at 30 m/pixel. This dataset is available in GEE as ‘TERN/AET/CMRSET_LANDSAT_V2_2’.

2.2.3. Vector Data

Two vector datasets were used as input in the GEE workflow: i) the Macquarie-Castlereagh River catchment boundary (polygon); ii) a set of landuse classification layers created for the entire state in 2004, 2013 and 2017. This layer was created based on *in situ* observations and high-resolution aerial photo data. For each crop season, the closest landuse layer was selected. For example, the 2012 crop year uses the 2013 landuse and 2019 crop year uses 2017 landuse.

2.3. Methodology

A 3-phase methodology was designed to map irrigated areas with the datasets at hand. Figure 2 presents a flowchart representation of each step together with infographics depicting the key intermediary products.

1. Phase I: creation of a seasonal maximum NDVI composite in GEE (Figure 3) using all images from 1 December to 31 March and excluding cloudy and rain-affected images.
2. Phase II: extraction of irrigated area polygons from maximum NDVI composite using a grid-based thresholding method and morphological image processing.
3. Phase III: manual validation of irrigated areas using ancillary datasets such as evapotranspiration, landuse and NDVI phenology.

2.3.1. Rainfall Filtering and Image Selection

Heavy to moderate rainfall can affect the greenness of vegetation on the images. This is one of the major obstacles for mapping irrigated areas, the green-flush vegetation is sometimes identified as irrigated and leads to a large number of false positives. Thus, a filter of 25mm rainfall within 5 days was applied to remove rainfall-impacted images from the seasonal maximum NDVI composite. This means that if it rains 25mm or more on a given day, any image captured within the next 5 days will be automatically discarded from the collection. The rest of the cloud-free (20% cloud threshold) images were used to compute maximum NDVI composite.

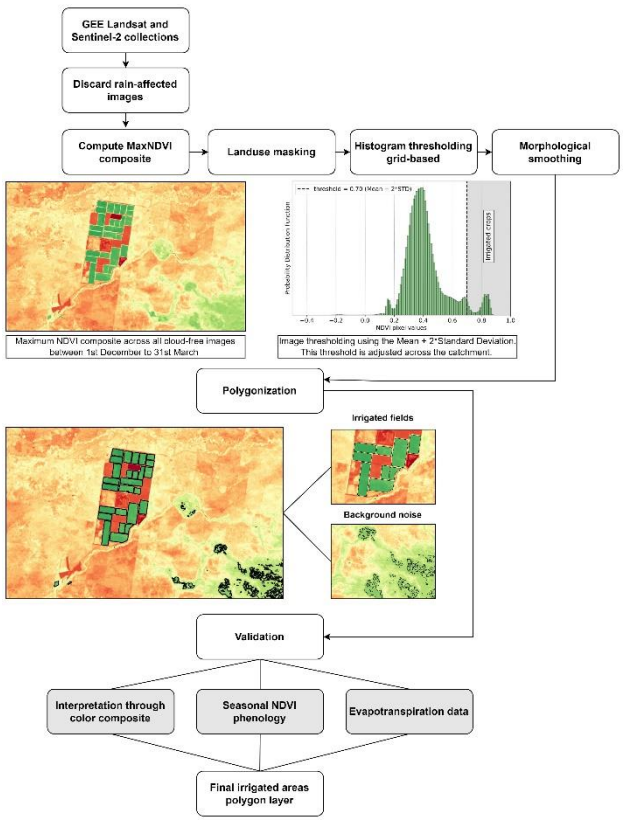

Figure 2. Flowchart of the methodology to map irrigated areas.

2.3.2. Maximum NDVI Composite

Time series of NDVI were essential for obtaining cropland phenological data and for distinguishing irrigated areas from rain-fed areas. Unlike other vegetation types, such as grassland and forest, croplands have unique observable characteristics in the stages of emergence, vegetative growth, senescence and harvest. These features mean a time-series of satellite images can be used to identify croplands and separate them from other vegetation types. NDVI (Figure 3) was derived using the following formula (Eqs. 1 and 2) based on the satellite sensor used.

$$NDVI = \frac{(NIR-Red)}{(NIR+Red)}, \quad (1)$$

Within GEE, a maximum NDVI composite (Eq. 2) was constructed using all the selected images in cropping season (1 December to 31 March):

$$NDVI_C = \max(NDVI_1, NDVI_2, \ldots\ldots\ldots NDVI_n), \quad (2)$$

where NDVIc represents maximum NDVI composite while n represents the number of images available for a season.

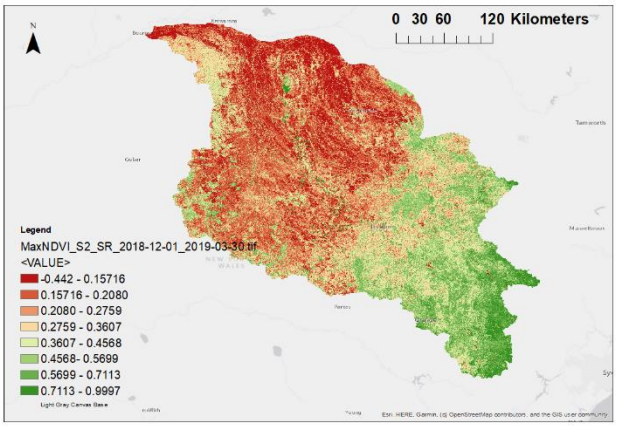

Figure 3. Maximum NDVI composite generated in GEE using all rain-filtered and cloud-free images for the 2018–2019 crop season.

2.3.3. Landuse Masking

The maximum NDVI composite was then combined with the NSW landuse vector data to ensure that only those areas that were potentially irrigated were retained. These areas were labelled as irrigated cropping categories in the NSW landuse layer. The following 8 classes were considered: ‘Grazing irrigated modified pastures’, ‘Grazing modified pastures’, ‘Irrigated cropping’, ‘Cropping’, ‘Seasonal horticulture’, ‘Perennial horticulture’, ‘Irrigated seasonal horticulture’ and ‘Irrigated perennial horticulture’. For more information on the Australian landuse classes, visit https://www.agriculture.gov.au/abares/aclump/land-use/alum-classification/alum-classes.

2.3.4. Thresholding

NDVI time series data derived and compiled using GEE from Landsat/Sentinel were then used to classify irrigated areas. Due to spatial variability in rainfall, topography, elevation and climate in the catchment, a local threshold was preferred to a global NDVI threshold. To account for those differences in the landscape, the catchment was separated in blocks and a separate

threshold was applied to each block based on the local histogram distribution of NDVI values. Figure 4a illustrates the maximum NDVI composite masked by landuse and divided into a 3x3 grid. Figure 4b shows the histogram of NDVI values within each block. It indicates there was strong variability between the lower (blocks 1–3) and higher parts of the catchment (block 7–9).

The NDVI threshold was calculated as the mean of the distribution plus 2 times the standard deviation (mean + 2*std). After applying the threshold, the resulting image was binary (Figure 5) – pixels with an NDVI above the threshold were considered irrigated crops. To avoid having ‘edge effects ‘for irrigated areas at the boundary between 2 adjacent blocks, a 4-km buffer was applied around each block and the lowest threshold among the overlapping buffer was used in the 4 km region. This ensured that individual paddocks were not partially detected because of a low NDVI threshold on one side and high NDVI threshold on the other.

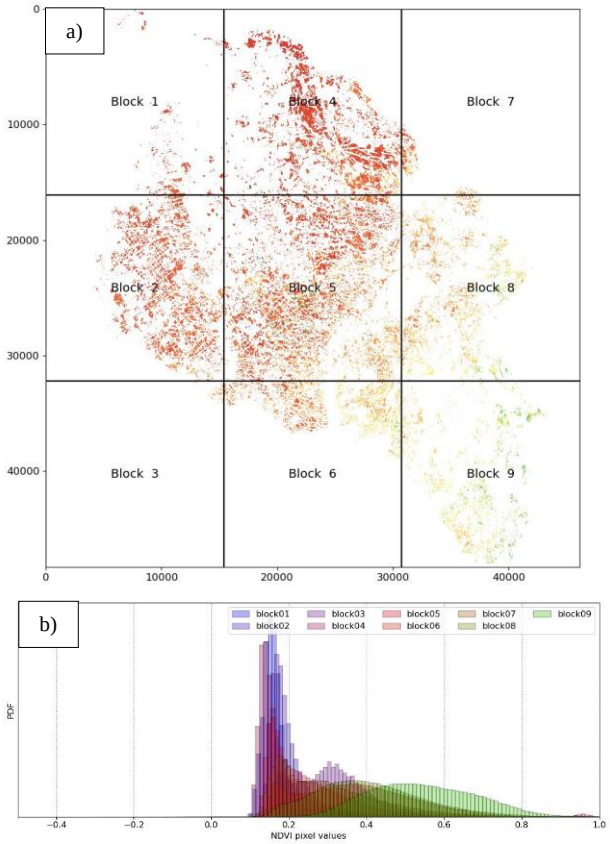

Figure 4. a) Maximum NDVI composite masked by landuse and divided in a 3X3 grid. b) Histogram distribution of NDVI values within each of the 9 blocks, showing variability in NDVI across the catchment.

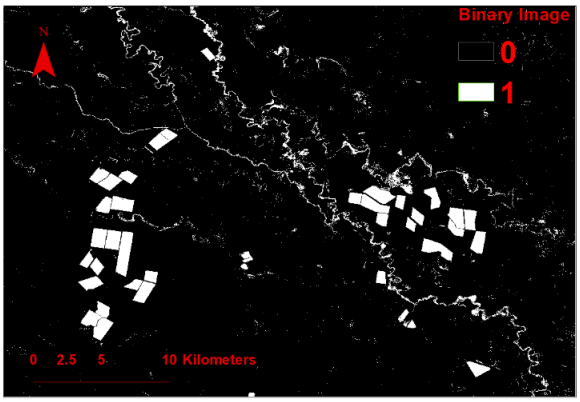

Figure 5. Binary image resulting from the application of a local NDVI threshold to separate the irrigated areas (white) from the non-irrigated areas (black).

2.3.5. Noise Removal by Morphological Image Processing

After applying the grid-based thresholding algorithm, the resulting binary image captured the irrigated paddocks but also contained much noise (Figure 5). This noise is created by pixels with high NDVI that are not irrigated paddocks (e.g., riparian vegetation).

To remove the noisy pixels, morphological operators were used to only retain large compact shapes that were likely to be irrigated paddocks. A binary opening with a 3x3 (Landsat) or 5x5 (Sentinel-2) square element was applied to the binary image. This consisted of an erosion, followed by a dilation. Therefore, it removed small and ‘skinny’ features while conserving compact shapes. Subsequently, another morphological filter was applied to only retain features with at least 11 (Landsat) or 100 (Sentinel-2) connected pixels, which corresponds to a minimum area of 1 hectare. Finally, a gap-filling operation was applied to fill holes inside the detected paddocks. All these operations were carried out using the *scikit-image* python library (https://scikit-image.org/docs/stable/api/skimage.morphology.html).

After these morphological operators were applied to the binary image, the image was vectorised to create individual polygons. The resulting polygons are indicated in Figure 6, with a backdrop of maximum NDVI composite. While the red polygons are the final extracted polygons, the blue polygons show all the riparian vegetation that was removed with the morphological operations.

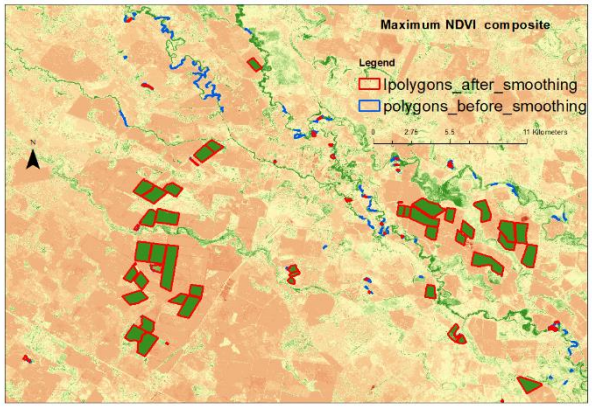

Figure 6. Resulting irrigated area polygons on maximum NDVI composite backdrop. The red polygons are the final extracted polygons, while the blue the blue polygons indicate the riparian vegetation that was removed by the morphological operations.

2.3.6. Validation

This final step checked the quality of the irrigated area maps produced. In it, the vector polygon layer generated in the automated workflow was validated manually and checked against 3 different sources of information:
(a) national map
(b) Seasonal NDVI phenology
(c) evapotranspiration data (ET)

a) **National Map**: In the national map, false-colour composites from Landsat and Sentinel-2 data could be visualised interactively. The polygon layer was overlayed over a backdrop of false-colour composite images using Green, SWIR and NIR bands, as shown in Figure 7. This visual interpretation helped identify 'false positives' – polygons that were identified as irrigated areas but which did not show the typical crop growth cycle across the season. https://nationalmap.gov.au/

(b) **Seasonal NDVI**: Besides false colour composite, NDVI value for individual images over the season were also checked to confirm the correct detection of irrigated area. Crop phenology is seasonal and time-dependent where NDVI value changes with the growth of the vegetation and drop before the harvesting as shown in figure 7. With this step, rainfall-affected polygons can be removed and others can be cross-checked to make sure their NDVI pattern follows crop growth. As discussed previously in the introduction, NDVI values change with growth and maturity of the planted crop and decreases again before the harvest. As an example, Figures 7a and 7b show the early stages of vegetative growth in December/January followed by the peak period of vegetative growth (Figures 7c and 7d) and senescence where the greenness drops (Figure 7 e).

(c) **Evapotranspiration Check and Confirmation**: ET data provided a further line of evidence to confirm that the mapped polygons were correctly identified as irrigated areas. The ET monthly data was used to check that the irrigated areas had ET values above 3.5 mm/day for at least for 2 consecutive months. If ET value was more than the threshold value, it confirmed that the paddock was wet and not at a water-stress level. After harvesting, if there is no crop, the paddock dries out and the ET value for that paddock also drops.

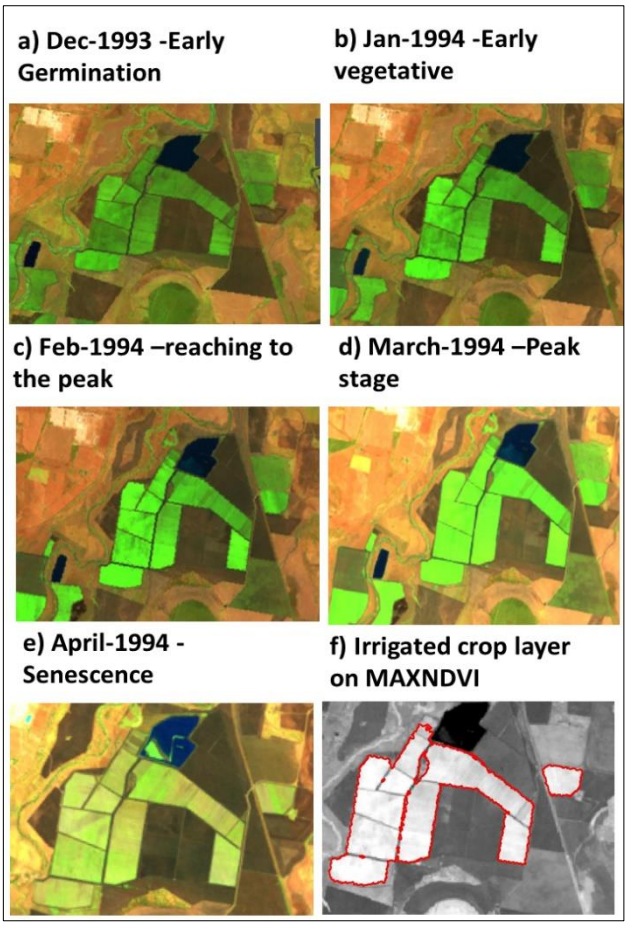

Figure 7. a–e: false-colour (Green, SWIR and NIR) imagery showing the vegetative growth from early germination (a), to crop green-up phase (b), reaching peak (c), peak (d) and senescence (e). f: maximum NDVI seasonal composite and validated irrigated crop areas.

These additional checks ensure that the final irrigated area map is of high quality and does not contain false positives.

3. Results and Discussion

This section presents a validation the satellite-derived irrigated areas. It then applies this dataset to calibrate catchment-scale hydrological models. Finally, it uses a long-term dataset of irrigated areas to investigate the relationship between irrigated areas, rainfall and climate patterns over the past 35 years.

3.1. Accuracy Assessment

The accuracy of the satellite-derived irrigated areas was assessed against 4 years of *in situ* surveys obtained from an Irrigator Behaviour Questionnaire (IBQ) in the northern Murray Darling Basin. Figure 8 shows the comparison between total irrigated areas derived from satellite and total irrigated areas reported in the questionnaire in the Border River catchment. The accuracy assessment indicates that the satellite-derived irrigated areas have a bias of ~1,750 ha and root-mean-squared-error (RMSE) of ~3,350 ha. This is a good agreement for the purposes of monitoring irrigated areas as it represents a 10% bias and ~20% RMSE.

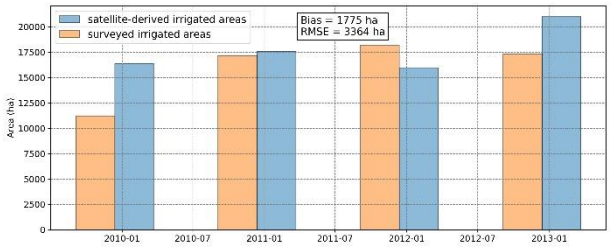

Figure 8. Comparison between satellite-derived irrigated areas and in situ irrigator questionnaire data for 4 years in the Border River catchment.

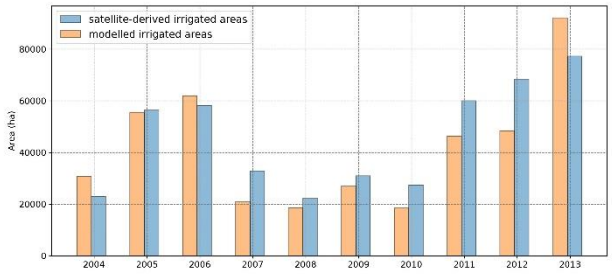

Figure 9. Satellite-derived irrigated areas in Gwydir catchment used for calibrating a catchment-scale hydrological model.

3.2. Calibration of Catchment-Scale Hydrological Models

Large-scale irrigated crop area time-series are crucial to calibrate catchment-scale hydrological models. These models are key to estimating current and future water usage and informing water planning strategies. To illustrate this point, a 10-year time-series of satellite-derived irrigated areas was used to calibrate a catchment-scale hydrological model in the Gwydir catchment in the northern Murray Darling Basin. Figure 9 shows the time-series of total irrigated areas derived from the satellite data and the corresponding irrigated areas predicted by the model. This shows how satellite-derived irrigated areas can be leveraged to calibrate the hydrological models and ensure that the models are correctly capturing irrigation processes across hundreds of properties over large catchments. The differences between the satellite observations and the hydrological model outputs are exacerbated during very wet years, as shown during 2011 and 2012 when the Gwydir catchment received above-average annual rainfall.

3.3. Long-term Patterns in Irrigated Areas and Links to Climate Drivers

By leveraging publicly available satellite imagery and the workflow described above, it is possible to monitor irrigated areas over large spatial and temporal scales. It is also possible to gather unique insights into spatial and temporal variability. For this paper, a dataset of total irrigated areas in the Macquarie-Castlereagh catchment (Figure 1) was compiled from 1987 to 2023. The time-series is shown in Figure 10a. It indicates strong inter-annual variability in total irrigated areas, with the minimum variability observed in 2019–2020 (10,000 ha of irrigated crops) and the maximum in 1996–1997 (67,000 ha).

One of the main drivers for the variability in irrigated areas is the availability of freshwater resources. Therefore, a corresponding time-series of total annual rainfall averaged over catchment was compiled with data from SILO (https://www.longpaddock.qld.gov.au/silo/). Additionally, the dominant climate driver affecting rainfall patterns in south-east Australia is the El Niño/Southern Oscillation (ENSO). The phase and intensity of ENSO is represented by the Multivariate ENSO Index (MEI), downloaded directly from NOAA (https://psl.noaa.gov/enso/mei/). Both these datasets are presented in Figure 10b and indicate a strong relationship, with dry periods associated with El Niño and wet periods with La Niña.

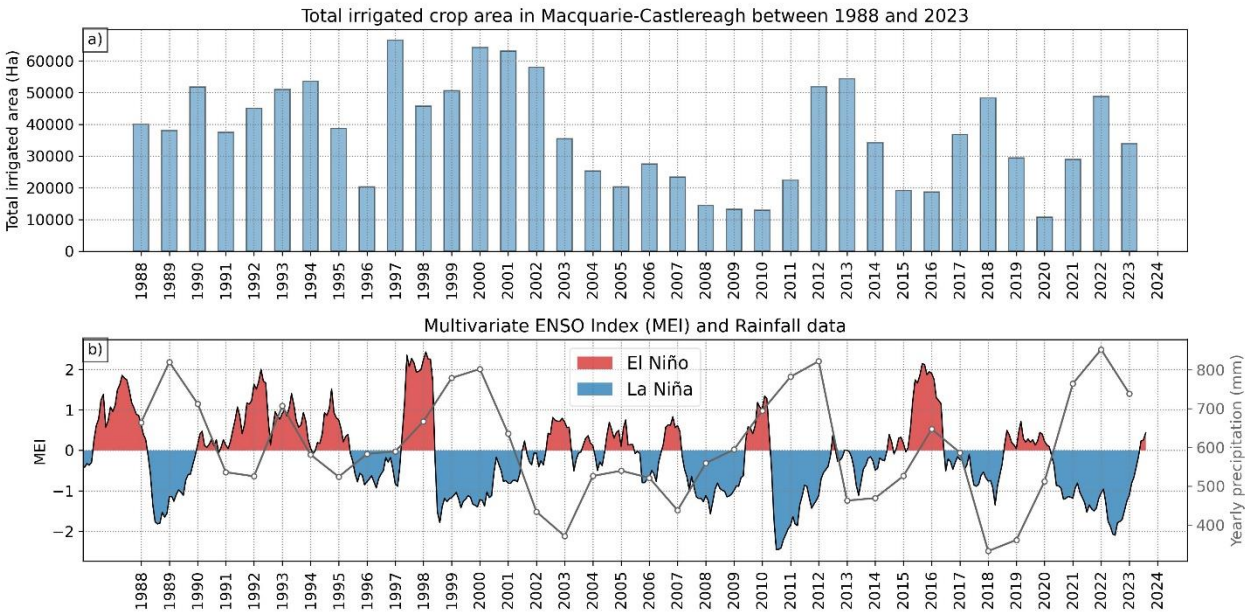

Figure 10. a) 35-year time -series of total irrigated area in Macquarie -Castlereagh catchment (1987-2023). b) Monthly time series of amplitude and phase of El Niño/Southern Oscillation and annual rainfall in the catchment.

The period between 1988 and 2000 was characterised by average rainfall and total irrigated areas remained relatively high. From 2001 to 2010, this region experienced the Millennium Drought, one of the worst droughts on record. During this period, very few areas in the catchment were irrigated – a total of only 10,000–20,000 ha, compared to 40,000–50,000 ha in the early 1990s. The strong La Niña phase of 2011–2012 brought high rainfall and total irrigated areas suddenly increased to the 50,000 ha mark. This was followed by an El Niño phase in 2016 that resulted in below-average rainfall and reduced total irrigated area. This drought last until 2020 and resulted in irrigated areas dropping to their lowest extent in 35 years. In 2020, rainfall returned with a ‘triple-dip’ La Niña. This resulted in an abundance of freshwater resources and increased total irrigated area (the highest in a decade).

This long-term time-series provides key insights into the availability of water for irrigation over the last 4 decades and the relationship between the climate and irrigation. This interannual variability is important for designing water management and water sharing strategies for different users within a catchment.

4. Concluding Remarks

This paper presents an operational methodology to extract irrigated crop areas from publicly available satellite imagery, namely Landsat and Sentinel-2. The 3-phase methodology combines NDVI imagery and ancillary environmental datasets (evapotranspiration data and a land use classification) with advanced image processing algorithms (local thresholding and morphological operations) in a cloud environment (Google Earth Engine) for robust, efficient and rapid processing. The automatically extracted irrigated area polygons are then quality-controlled with a manual validation. As part of this process, the satellite-derived irrigated areas used in this paper were compared against *in situ* data obtained from irrigators questionnaires and indicated a good agreement between the two datasets.

The developed methodology was then employed to create 35-year time-series of irrigated areas over an entire catchment in south-east Australia. The long-term dataset reveals interesting temporal patterns in the irrigation of cultivated areas. Strong inter-annual variability is observed that is linked to rainfall patterns and ENSO phases.

In conclusion, satellite remote sensing can provide essential information on irrigation practices to support water management strategies. This paper illustrates the value that satellite-derived observations of irrigated areas could have for water management purposes. Future work should focus on improving the methodology by using additional datasets such as higher resolution optical imagery like Planet (3 m/pixel), hyperspectral imagery and Synthetic Aperture Radar products.

Acknowledgments

Funding for this work was provided by the Australian Government under the Hydrometric and Remote Sensing (HNRS) project. The authors acknowledge the contribution of Meredith Tucker for editing the manuscript. We also thank Ammar Mahmood from NGIS who helped migrate the methodology in Google Earth Engine. The Evapotranspiration data is provided by CSIRO (CMRSET product). The rainfall data is provided by Queensland Government (SILO). The landuse classification layers are provided by the NSW Land Use Program (DCCEEW).

References

Abdi, A.M. (2020) ‘Land Cover and Land Use Classification Performance of Machine Learning Algorithms in a Boreal landscape using Sentinel-2 Data’, *GI Science & Remote Sensing*, 57:1–20.

Jensen, J.R. (2021): *Introductory Digital Image Processing: A Remote Sensing Perspective*, 4th Ed, Pearson.

Magidi, J.; Nhamo, L.; Mpandeli, S. & Mabhaudhi, T. (2021) ‘Application of the Random Forest Classifier to Map Irrigated Areas Using Google Earth Engine’, *Remote sensing*, 13:876. 10.3390/rs13050876.

Rasel, S. M. M.; Chang, H. C.; Ralph, T. J.; Saintilan, N. & Diti, I. J. (2021) ‘Application of feature selection methods and machine learning algorithms for saltmarsh biomass estimation using Worldview-2 imagery’, *Geocarto International*, 36(10): 1075–1099. 10.1080/10106049.2019.1624988.

Rouse, J. W.; Haas, R. H.; Schell, J. A. & Deering, D. W. (1973) ‘Monitoring Vegetation Systems in the Great Plains with ERTS (Earth Resources Technology Satellite)’ *Proceedings of 3rd Earth Resources Technology Satellite Symposium*, Greenbelt, SP–351:309–317.

Teleguntla, P.; Thekabail, P.; Oliphant, A.; Xiong, J.; Gumma M. K.; Conglaton, R. G. & Yadav, K Huete.; A. A. (2018) ‘A 30-m Landsat derived cropland extent product of Australia and China using Random Forest machine learning algorithms on Google Earth Engine cloud computing platform’, *ISPRS Journal of Photogrammetry and Remote Sensing*, 144:325–340, 10.1016/j.isprsjprs.2018.07.017.

Van Loon, A. F.; Stahl, K.; Di Baldassarre, G.; Clark, J.; Rangecroft, S.; Wanders, N.; Gleeson, T.; Van Dijk, A. I. J. M.; Tallaksen, L. M.; Hannaford, J.; Uijlenhoet, R.; Teuling, A. J.; Hannah, D. M.; Sheffield, J.; Svoboda, M.; Verbeiren, B.; Wagener, T. & Van Lanen, H. A. J. (2016) ‘Drought in a human-modified world: Reframing drought definitions, understanding, and analysis approaches’, *Hydrology and Earth System Sciences*, 20:3631–3650, 10.5194/hess-20-3631-2016.

Wardlow, B .D. & Egbert, S. L (2008) ‘Large-area crop mapping using time-series MODIS 250 m NDVI data: An assessment for the U.S. Central Great Plains’, *Remote Sensing of Environment*, 112:1096–1116, 10.1016/j.rse.2007.07.019.