

SHAZAM: Self-Supervised Change Monitoring for Hazard Detection and Mapping

Samuel Garske , *Graduate Student Member, IEEE*, Konrad Heidler , Bradley Evans , Kee Choon Wong ,
and Xiao Xiang Zhu , *Fellow, IEEE*

Abstract—The increasing frequency of hazardous events due to climate change underscores an urgent need for effective hazard monitoring systems. Current approaches rely on expensive labelled datasets, struggle to distinguish hazardous changes from seasonal changes, or require multiple observations for to confirm when a potential hazardous event has occurred. To address these challenges, this work presents Self-Supervised Change Monitoring for Hazard Detection and Mapping—SHAZAM. SHAZAM uses a lightweight conditional UNet model to generate images of a region of interest (ROI) for any day of the year, allowing for the direct modelling of seasonal changes. A modified structural similarity measure compares the generated images with real satellite observations to compute both image-level anomaly scores, and pixel-level hazard maps. In addition, a seasonal threshold is introduced to further reduce the need for dataset-specific optimisation. SHAZAM was evaluated on four diverse datasets that contain wildfires, burned regions, extreme and out-of-season snowfall, floods, droughts, algal blooms, and deforestation, and achieved F1 score improvements of between 0.066 and 0.234 over existing methods. This was achieved primarily through a higher recall over other models, while using only 473 K parameters. SHAZAM demonstrates superior mapping

capabilities through a higher spatial resolution and the suppression of background features for both immediate and gradual hazards. SHAZAM offers an effective and generalisable solution for hazard detection and mapping across different geographical regions and a diverse range of hazards, and highlights the potential of further self-supervised method contributions for hazard detection and mapping.

Index Terms—Anomaly detection, change detection, deep learning, hazard detection, satellite image time series (SITS), self-supervised learning.

I. INTRODUCTION

ENVIRONMENTAL hazards are any natural or human-induced phenomenon that has the potential to cause harm or damage to communities, ecosystems, or the environment as a whole. Common examples include floods, storms, landslides, oil spills, earthquakes, and wildfires. Methods that automatically detect hazards allow organizations to actively monitor a region of interest (ROI), and mapping them provides crucial geographic information needed to respond to and manage their impact [1]. ROIs can include areas prone to hazards, regions that require environmental protection, or remote areas that lack ground-based monitoring infrastructure. Monitoring and managing environmental hazards is of increasing importance as hazardous events are becoming more frequent in a changing climate [2]. This urgency is highlighted by the United Nations’ “Early Warnings for All” initiative, which aims to achieve global coverage by 2027, addressing the critical gap where one-third of the world’s population lacks access to adequate warning systems for natural hazards [3].

Satellites with onboard cameras are an effective tool for remotely monitoring an ROI, given their regular coverage of much of the Earth’s surface. Regularly collecting images of the same ROI over time creates a satellite image time series (SITS), providing a temporal record of the landscape that can be used to identify unexpected changes and hazards (see Fig. 1). The large and growing amount of available global historical SITS data provides unprecedented land monitoring capabilities, with missions such as Landsat and Sentinel-2 capturing high-resolution multispectral images [4], [5]. Synthetic aperture radar (SAR) satellites such as Sentinel-1 are also a powerful choice for hazard monitoring because they can penetrate cloud cover and capture images at night [6]. Various methods for hazard detection and mapping with these satellite systems cover floods, landslides, deforestation, wildfires, and algal blooms [6],

Received 28 February 2025; revised 26 April 2025; accepted 10 June 2025. Date of publication 12 June 2025; date of current version 1 July 2025. This work was supported in part by the German Federal Ministry of Education And Research (BMBF) through the international future AI lab “AI4EO – Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond,” under Grant 01DD20001, and in part by the Australian Research Council through the Industrial Transformation Training Centre, Australian Research Council (ARC) Training Centre for CubeSats, UAVs and Their Applications (CUAVA), under Grant IC170100023. The work of Samuel Garske was supported by international Future AI lab “AI4EO – Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond through Beyond Fellowship at TUM. The work of Konrad Heidler was supported by the German Federal Ministry for Economic Affairs and Climate Action in the framework of the “National Center of Excellence ML4Earth” under Grant 50EE2201C and in part by the Munich Center for Machine Learning. The work of Xiao Xiang Zhu was supported by the Munich Center for Machine Learning. (*Corresponding author: Samuel Garske.*)

Samuel Garske and Kee Choon Wong are with the School of Aerospace, Mechanical, and Mechatronic Engineering, The University of Sydney, Sydney, NSW 2006, Australia, and also with the Australian Research Council (ARC) Training Centre for CubeSats, UAVs and Their Applications (CUAVA), Australia (e-mail: sam.garske@sydney.edu.au).

Konrad Heidler was with the SiPEO, Technical University of Munich, 80333 Munich, Germany. He is now with Hula Earth, 80335 Munich, Germany.

Bradley Evans is with the School of Environmental and Rural Science, The University of New England, Armidale, NSW 2350, Australia, and also with the Australian Research Council (ARC) Training Centre for CubeSats, UAVs and Their Applications (CUAVA), Australia.

Xiao Xiang Zhu is with the Chair of Data Science in Earth Observation (SiPEO), Department of Aerospace and Geodesy, School of Engineering and Design, Technical University of Munich, 80333 Munich, Germany, and also with the Munich Center for Machine Learning, Technical University of Munich, 80333 Munich, Germany.

Data is available online at <https://github.com/WiseGamgee/SHAZAM>.
Digital Object Identifier 10.1109/JSTARS.2025.3579330




Fig. 1. This Sentinel-2 SITS example shows a wildfire event in Sequoia National Park, California, progressing from normal conditions (left), to active wildfires (middle), and to burned regions (right). The corresponding hazard detection heatmaps were created by the proposed method, SHAZAM.

[7], [8], [9], [10]. However, the sheer volume and complexity of SITS data have driven the development of automated computer vision and machine learning approaches for this task.

Change detection has emerged as a fundamental approach for automated hazard monitoring using satellite imagery. Change detection is the process of identifying differences in the state of an object or phenomenon by observing it at different times [11]. In the context of remote sensing, change detection methods typically map the different pixels between two separate images of the same location. This approach is particularly relevant for hazard mapping, as models can learn to identify changes between satellite images that directly indicate a hazard.

However, current change detection approaches face several key limitations. Supervised methods require extensive labeled datasets, which are expensive and time consuming to create [12], [13], particularly for the diverse range of potential hazards across different geographical locations. Unsupervised approaches eliminate the need for labels, but struggle to distinguish hazards from normal seasonal changes and are better suited to postevent mapping rather than active monitoring. Methods that attempt continuous monitoring, such as the continuous monitoring of land disturbance (COLD) method [14] and near-real-time monitoring of land disturbance based on time-series of harmonized reflectance (NRT-MONITOR) [15], address seasonal variations but are computationally intensive and require multiple posthazard observations, hindering timely detection.

Recent advances in self-supervised learning and anomaly detection offer promising directions for addressing these challenges [16]. These approaches can learn normal patterns from unlabeled data and identify deviations that may indicate hazards, as demonstrated in video anomaly detection [17], [18]. However, adapting these methods to satellite imagery presents unique challenges, particularly in handling the irregular temporal sampling of SITS data and modeling seasonal dynamics throughout the year. To address these challenges, this work makes three following key contributions.

- 1) SHAZAM—a self-supervised change monitoring method for hazard detection and mapping. SHAZAM is a novel application of self-supervised learning for hazard monitoring, allowing the detection and mapping of potential hazards without requiring labeled examples or multiple confirmatory observations.
- 2) Three technical modules.
 - The seasonally integrated Unet (SIU-Net): a lightweight convolutional architecture that translates a baseline ROI image into its expected appearance for any day of the year;
 - A structural difference index measure (SDIM): A slightly modified SSIM that yields a region-level anomaly score and a pixel-level anomaly heatmap, enabling higher resolution hazard mapping than existing unsupervised methods;
 - A novel seasonal threshold for identifying if an image’s anomaly score is sufficiently high to flag a potential hazard.
- 3) Extensive validation—experiments on four datasets covering diverse geographical regions and seven hazard types (wildfires, burned regions, extreme and out-of-season snowfall, algal blooms, drought, deforestation, floods) show that SHAZAM outperforms state-of-the-art unsupervised hazard monitoring methods in both detection accuracy and mapping resolution.

The rest of this article is organized as follows. Section II reviews related work on change detection and hazard mapping. Section III describes the proposed methodology, SHAZAM. Section V presents the experimental results, analysis, and highlights the core areas of future work. Finally, Section VI concludes this article.

II. RELATED WORK

Traditionally, change detection approaches encompassed algebra-based methods, transformation, classification, advanced models, GIS, and visual analysis [19]. More recently, deep learning has become the dominant approach for supervised change detection due to its ability to learn complex spatial and spectral representations from satellite imagery [20], [21], [22]. Supervised approaches span various architectures, such as encoder-decoder networks [23], [24], [25], Siamese networks [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], [38], generative adversarial networks [39], [40], [41], recurrent neural networks [42], [43], [44], [45], [46], [47], and lightweight architectures [48], [49], [50].

However, supervised change detection methods face three key limitations for hazard monitoring. First, they typically rely on carefully selected bitemporal pair of images from before and after a hazard has occurred. These approaches are better suited for postevent mapping than active monitoring and detection. Second, they lack mechanisms to handle seasonal variations, which means that they require images from similar seasons to avoid mistaking normal seasonal changes for hazards. Third, and most critically, they require extensive labeled datasets, which are expensive and time-consuming to create for the diverse range of potential hazards across different geographical locations [12], [13]. In many cases, such methods are only demonstrated on a single ROI and hazard type. These limitations have motivated the development of unsupervised approaches—i.e., those that do not use labeled datasets—which is the core focus of this work.

A. Unsupervised Change Detection: Towards Continuous Monitoring

Unsupervised methods identify changes between images without needing labeled change maps for training, eliminating the resource-intensive process of manual labeling. Traditional bitemporal approaches use alternative learning objectives, such as the stacked sparse autoencoder proposed by [51]. This autoencoder reconstructs pixels from two sensor images, classifying changes based on reconstruction errors, the difference between model predictions and actual values, with changed pixels typically showing larger errors. Although avoiding labeled samples, it still depends on carefully chosen images from similar seasons, restricting continuous monitoring applications. Moving toward multitemporal SITS methods, [52] applied an LSTM framework using multiple SAR images to map flood events by reconstructing shuffled pixel time series sequences. Though representing progress in leveraging SITS, this approach lacked seasonal modeling and required selecting unaffected training regions retrospectively, making it more suitable for postevent mapping than monitoring.

Some methods now explicitly address seasonal variations in SITS. UTRNet [53] embeds multiyear temporal information within a convolutional LSTM network, using both seasonal cycle position and temporal distance data with self-generated labels to differentiate genuine changes from seasonal variations. However, it requires costly retraining with all historical data for new observations, and its detection capacity is limited by its underlying prechange model. Similarly, UCDFormer [54] employs a lightweight transformer with domain adaptation, attempting to suppress seasonal effects through an affinity-weighted translation loss that measures both spectral distribution distances and structural content differences between images. This suppresses changes caused by seasonality, but still requires image selection near identified events rather than identifying the events themselves.

More recent work has focused on lightweight solutions suited to monitoring. RaVAEn [55] implements a variational autoencoder for onboard change event detection, compressing Sentinel-2 image patches to detect and map multiple hazards including floods, fires, landslides, and hurricanes. CLVAE [56] advances

this using contrastive learning for SAR flood mapping, and similarly handling irregular capture times without labeled data. While these operational methods demonstrate progress for practical monitoring, they still lack robust seasonal integration to distinguish real changes from seasonal variations.

B. Seasonal Modeling Through Disturbance Detection

An area that provides insight into modeling seasonal dynamics when detecting potential hazards is disturbance detection, which can be considered a subfield of unsupervised change detection. Disturbance detection has traditionally focused on detecting abrupt events that result in the sustained disruption of an ecosystem, community, or population [57], [58]. Methods such as the breaks for additive seasonal and trend algorithm (BFAST) [59] and the landsat-based detection of trends in disturbance and recovery algorithm (LandTrendR) [60] model seasonal variations and long-term trends of a region using sequential multispectral satellite images, and can identify abrupt changes caused by natural or human-induced hazards.

BFAST is a more statistical approach that identifies disturbances using breakpoints, which are times when the SITS-derived normalized difference vegetation index (NDVI) significantly alters from the modeled trend after the removal of seasonality. BFAST uses piecewise linear regression to model the NDVI trend and employs dummy variables to represent annual seasonality cycles. LandTrendR employs a geometric method to identify breakpoints in the NDVI and in the normalized burn ratio (NBR) time series, using piecewise linear segments and angles between them to model trends. However, it does not directly model seasonality.

Disturbance-based methods have since been demonstrated for various applications, such as the detection and mapping of droughts, urbanization, forest degradation and deforestation, land temperature anomalies, grassland mowing, impacts on coastal wetlands, and burned regions [61], [62], [63], [64], [65], [66], [67], [68]. However, these approaches tend to use specific spectral indices and rule-based methods to detect disturbances in a particular biome. This limits their ability to detect and map multiple hazards across diverse geographic regions.

More generalized methods that use multiple spectral bands were developed for pixelwise disturbance detection and monitoring such as the COLD method [14], and NRT-MONITOR [15]. COLD uses a limited Fourier series to model intra-annual seasonality and interannual trends in Landsat images, predicting each band's reflectance values. The Fourier series models seasonal cycles over yearly, half-yearly, and four-month intervals. Deviations deemed statistically significant are identified as disturbances. NRT-MONITOR employs a similar Fourier series for seasonality on harmonized Landsat-Sentinel images but uses a forgetting factor for near real-time updates. This eliminates the requirement for storing complete historical time series of each pixel to identify potential hazards. It operates faster than COLD and flags disturbances with a shorter lag-time.

Despite efforts to develop faster near real-time methods [14], [15], [63], [69], disturbance detection methods are extremely computationally demanding and are typically implemented on

high-performance computers (HPCs). They also require multiple posthazard observations to confirm a disturbance. This reduces false disturbance alarms, but comes at the cost of delayed hazard detection. This hinders the reaction of disaster and hazard management agencies, where a fast response time is critical. Pixelwise approaches also lack the spatial context required for effective high-resolution image analysis [70]. Few disturbance detection studies have investigated generalized spatially aware methods [69], and none involve a framework that simultaneously flags multiple hazards at a regional level while mapping them at the pixel level.

C. Self-Supervised Learning and Anomaly Detection

Although disturbance detection successfully models seasonality, self-supervised learning offers a more data-driven and efficient approach to change monitoring. Self-supervised learning is a subset of unsupervised learning in which a model is trained on a traditionally supervised task, such as regression or classification, without human-annotated labels [16], [71], [72]. Self-supervised learning often involves the use of machine-generated labels, contrastive learning, which separates different observations while clustering similar ones, or autoencoding, where models learn from reconstructing the input [16]. Self-supervised learning has become popular for creating foundation models on large unlabeled datasets in remote sensing [73], [74], [75], which are later fine-tuned for downstream tasks with smaller and more niche datasets.

Self-supervised anomaly detection methods for videos offer a unique perspective toward change monitoring for an ROI, given their ability to identify unexpected events within time-series images. Anomaly detection methods are typically trained using only normal video footage, gaining a strong understanding of the expected dynamics within a given scene [76], [77], [78], [79], [80], [81]. This could include normal traffic patterns for a certain intersection or normal human behavior in security camera footage. During testing, anomalous frames are flagged when the model struggles to represent them. Yu et al. [82] even demonstrated that the use of contaminated training data (that is, training data containing anomalous images) still performed well. This was based on the assumption that anomalies are rare occurrences and deep learning networks will bias themselves toward representing the majority of normal images. This is outlined as an important finding, as ensuring high-quality annotation is a difficult task in remote sensing [12], [13].

In the realm of remote sensing, video-based anomaly detection methods have been adapted for aerial videos captured by drones. Jin et al. [17] proposed ANDT, a self-supervised hybrid transformer-convolutional network. ANDT uses six consecutive images to predict the subsequent image, using the prediction error to identify unexpected events. Images with prediction errors one standard deviation above the mean training error are classified as anomalies. The transformer-based encoder improved time-series feature mapping and future frame prediction, outperforming other well-known anomaly detection methods such as MemAE [83] and GANomaly [84]. Tran et al. [18] enhanced the approach by adding temporal cross-attention to the

encoder outputs, improving spatial and temporal feature capture and surpassing ANDT's performance.

Although video-based detection methods offer promising directions, their direct application to SITS faces one major challenge. Unlike the uniform temporal sampling of video data, SITS typically presents in irregular intervals due to cloud cover and other environmental factors. This irregularity particularly affects transformer architectures, which rely on learning from dense, regular sequences to model temporal dependencies [85], [86].

Recent self-supervised transformer architectures for SITS have incorporated temporal encoding mechanisms, allowing models to handle the irregular and inconsistent timing between satellite image acquisitions [87], [88], [89]. These models employ input masking or corruption techniques during pretraining to learn spatial, spectral, and temporal features before tackling downstream supervised classification and segmentation tasks. For example, SITS-Bert applies noise to random pixels in a time-series and trains the network by denoising them. SITS-Former masks a random patch (5×5 pixels) in a patch time-series and learns by reconstructing it. U-BARN first encodes time-series patches using a CNN, then randomly masks some patch features, and finally reconstructs them to train the network.

A critical challenge in these approaches is to effectively encode seasonality. These architectures encode the day of the year using positional sine-cosine embeddings, borrowed from the original transformer architecture [90]. However, since these embeddings were originally designed to encode sequential positions rather than circular temporal patterns, they may not fully capture the cyclical nature of seasonal changes. An alternative approach, used by the pretrained remote sensing transformer, converts the month of image acquisition into cyclical coordinates [91]. This offers a more natural and direct embedding of seasonal patterns into the model, and is similar to the temporal modeling used in disturbance detection. It is clear that opportunities remain to develop more sophisticated seasonal embeddings specifically for anomaly detection.

Overall, the integration of self-supervised learning principles for SITS and video anomaly detection presents a promising direction for change monitoring. By modeling the normal spatio-temporal behavior of an ROI across different seasons, self-supervised approaches could learn to distinguish genuine changes from regular seasonal variations without the need for labeled data. Directly embedding temporal information not only helps learn seasonal behavior, but it offers a combined solution to address the challenge of irregularly spaced SITS.

D. Summary

This section's review of hazard detection and mapping approaches reveals four main directions: supervised change detection, unsupervised change detection, disturbance detection, and self-supervised video anomaly detection. Although supervised change detection methods provide a solid baseline for hazard mapping, they rely on costly and time-intensive human-annotated datasets and are better suited for postevent mapping

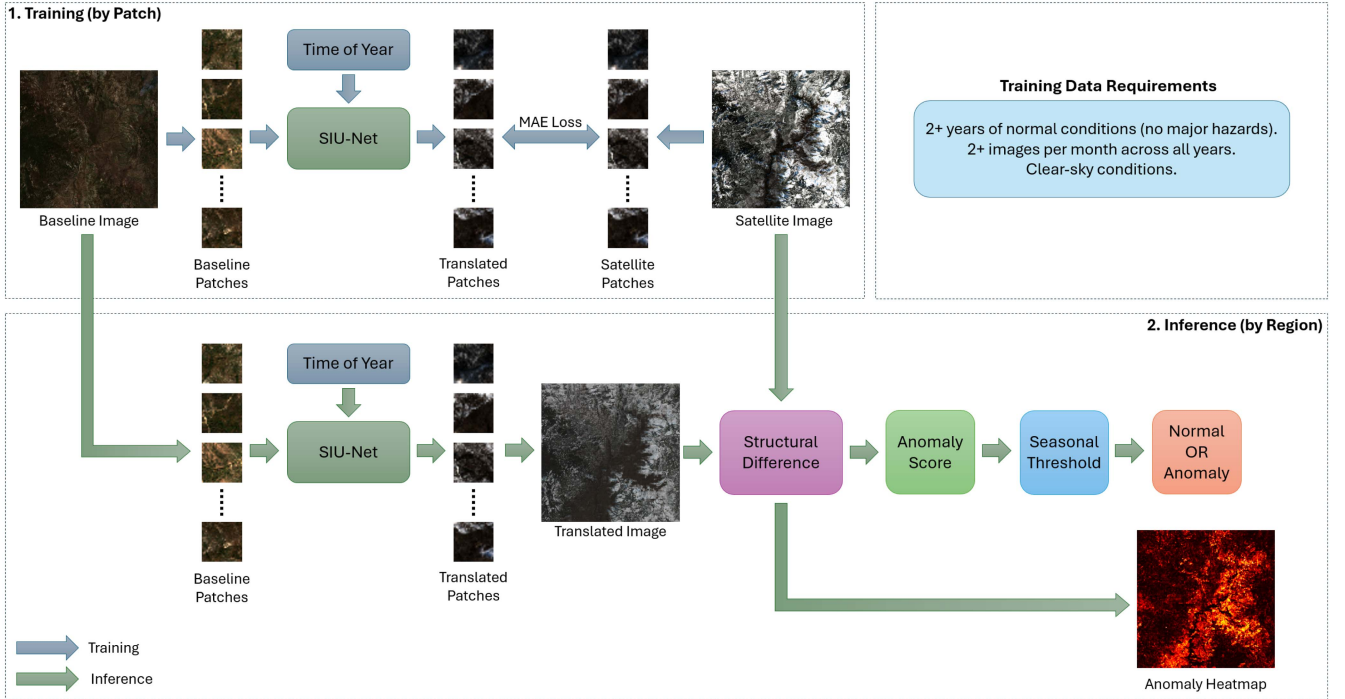


Fig. 2. SHAZAM—An overview of the proposed method. The top left visualizes the training stage, the bottom the inference stage for monitoring an ROI, and the top right highlights training data requirements.

than continuous monitoring. Unsupervised change detection approaches eliminate the need for labeled data and offer flexibility for detecting various types of hazards, but struggle to handle seasonal changes. Disturbance detection methods filter out seasonal changes but are computationally intensive and require multiple posthazard observations, hindering timely detection. Self-supervised methods combined with video-based anomaly detection show promise for ROI-specific hazard monitoring, though challenges remain to handle the irregular capture times of SITS and model seasonal patterns. In the next section, we introduce SHAZAM to address these limitations.

III. PROPOSED METHOD

A. SHAZAM Overview

This section introduces SHAZAM, a self-supervised change monitoring method for hazard detection and mapping (see Fig. 2). By using 2+ years of satellite images without any major hazards present, SHAZAM learns the normal seasonal patterns of a ROI. Using self-supervised learning, the SIU-Net model learns to translate a baseline image to any given satellite image conditioned by the day of the year they were captured. It does this by ingesting 32×32 pixel patches, rather than ingesting the whole image at once. After training, SIU-Net can effectively generate an image of the ROI for any day of the year. The baseline image is created by computing the median image (pixelwise) of all satellite images from the training dataset, excluding the validation patches.

In the inference phase (used for monitoring), SIU-Net generates an image of the ROI-based on the same day of the year

when a new satellite image is captured. A structural difference measure compares the satellite and SIU-Net images, producing an anomaly heatmap of unexpected pixel changes, serving as a hazard map. The average structural difference across all pixels in the anomaly heatmap gives the anomaly score, which is then used to detect a hazard. Due to seasonal variability in SIU-Net's translation performance, a seasonal error threshold is estimated from the training dataset and subtracted from the anomaly score. Final anomaly scores above zero indicate potential hazards, while those below zero are normal.

B. Preprocessing

All 32×32 patches (both inputs and outputs) are normalized using a min-max scaling approach based on percentile values. The values of the 1st percentile (x_1) and 99th percentile (x_{99}) values are calculated using all patches in the training dataset (across all input channels). These same percentile values are then used to normalize the validation and test patches using the following:

$$x_{\text{norm}} = \frac{x - x_1}{x_{99} - x_1}, \quad x \in \mathbb{R} \quad (1)$$

where x represents the original reflectance value, x_1 and x_{99} are the 1st and 99th percentiles of the training data, respectively, and x_{norm} is the normalized value. After normalization, values are clamped to the range $[0, 1]$ as follows:

$$x_{\text{final}} = \max(0, \min(1, x_{\text{norm}})). \quad (2)$$

This robust normalization approach effectively handles the heavy-tailed distributions typically present in satellite imagery,




Fig. 3. SIU-Net model architecture—the backbone of SHAZAM.




Fig. 4. Network blocks.

where extreme outliers (such as clouds, shadows, or strong reflections) could otherwise skew the scaling [92]. Using percentile-based normalization instead of absolute minimum and maximum values helps preserve the meaningful variation in the data while reducing the impact of these outliers.

C. Model Architecture

SHAZAM uses a conditional UNet-style model to translate 32×32 pixel baseline patches to their expected appearance for any day of the year, with an assumed 10 spectral channels for this work. This model is called SIU-Net, and the architecture is shown in Figs. 3 and 4. UNet is an effective architecture for change detection because it preserves the spatial structure with skip connections while capturing features from multiple scales [93]. These two qualities are essential for providing high-quality, high-resolution outputs that accurately represent the ROI.

The baseline patch is the primary input of SIU-Net, which is the median patch calculated from all patches in the training dataset (excluding validation patches). Each baseline patch provides the core spatial features of the region to SIU-Net (e.g., forests, rivers, lakes, mountains), serving as a reference point for temporal translation. This allows SIU-Net to act as a

conditioned generative model, enabling it to focus on seasonal transformation rather than memorizing spatial features. It also allows for a much smaller model compared to fully-generative alternatives (e.g., cVAE, cGAN, stable diffusion), which require significantly more parameters and data to accurately generate an ROI from scratch. The baseline patches are joined to form the baseline image of the entire ROI.

To perform seasonal translation efficiently, SIU-Net uses a lightweight architecture that integrates temporal information through time-position encodings concatenated as input channels. The network consists of a series of blocks (see Fig. 4); one double convolutional block ($2 \times \text{Conv2D}$ with GELU), two down-blocks (MaxPool + double convolutional block), and two up-blocks (bilinear upsampling + double convolutional block), with skip connections between feature spaces. All double convolutional block layers use 3×3 kernels with a stride of 1 and a padding of 1. The MaxPool layers use 2×2 kernels and a stride of 2. Feature channels progressively expand from 32, to 64, to 128 through the down-blocks before returning to the $32 \times 32 \times 10$ input dimensions via the up-blocks and a final Conv2D projection layer. The final Conv2D layer uses 1×1 kernels with a stride of 1 to reduce the concatenated feature maps, residuals, and time-position encodings to 10 channels. Skip connections play a dual role by concatenating both time-position encodings and residual features after each up-block, reinforcing temporal context while preserving the key spatial features of the patch during translation. SIU-Net uses the L1 loss function during training, also known as the mean absolute error (MAE), which calculates the absolute difference between each translated patch and the true values.

D. Season and Position Encoding

Capturing seasonal changes is one of the primary research gaps for change monitoring approaches. For SHAZAM, seasonality is injected into SIU-Net by converting the day of the year that an image was captured to a cyclical representation as




Fig. 5. Visualizing the seasonal encodings for the day of the year, t .

follows:

$$t_{\sin} = \sin\left(\frac{2\pi t}{365}\right) \quad (3)$$

$$t_{\cos} = \cos\left(\frac{2\pi t}{365}\right) \quad (4)$$

where t is the current day of the year. These encodings directly capture the cyclical nature of seasonality, inferring that a day in late December (e.g., day 350) is near early January (day 5). This is illustrated in Fig. 5. Another feature of these encodings is that they enable SIU-Net to interpolate between days that do not have corresponding images in the training set. This is crucial for handling irregular capture times with only a few years of historical SITS. Overall, these seasonal encodings allow SHAZAM to model both seasonal changes while handling irregular satellite observations.

The position encodings denote the location of each patch within the ROI image using normalized row and column coordinates as follows:

$$p_{\text{row}} = \frac{\text{row}}{n_{\text{patches}}} \quad (5)$$

$$p_{\text{col}} = \frac{\text{col}}{n_{\text{patches}}} \quad (6)$$

where n_{patches} is the number of patches per dimension, assuming a square image. Position encoding allows SIU-Net to better model seasonal dynamics on a local level, enhancing image translation for each individual patch. This relative coordinate system provides a foundation for potential model extensions to handle multiple regions through geographic conditioning. Both the season and position encodings are concatenated as extra channels to the input and to the residuals after each SIU-Net

convolutional block. As there are four encoding values (two seasonal and two position) there are four extra channels, and each value is duplicated across its respective channel (for example, a row encoding of 0.5 has all channel values set to 0.5).

E. Anomaly Map and Scoring

Unlike change detection where the main evaluation metric is pixelwise accuracy of the change map, SHAZAM's shift toward change monitoring focuses on imagewise event detection. Hazard mapping is a secondary objective. As such, SHAZAM treats hazard detection as a time-series anomaly detection problem, assigning each satellite image an anomaly score, which can be used to detect significant changes in the ROI. To do this, SHAZAM uses the structural difference between the real satellite image and SIU-Net's generate image for the same day of the year. The structural difference measure is based on the structural similarity index measure (SSIM [94]).

The SSIM compares the similarity of two images—in this case, the image generated by SIU-Net $\hat{y}$ and the real satellite image y —using three key components: luminance, contrast, and structure. For each pixel, these components are calculated in a local window (typically 11×11 pixels) that moves across the image in a sliding window fashion. Mathematically, the SSIM score of each pixel is defined as follows:

$$\text{SSIM}(\hat{y}, y) = \frac{(2\mu_{\hat{y}}\mu_y + k_1)(2\sigma_{\hat{y}y} + k_2)}{(\mu_{\hat{y}}^2 + \mu_y^2 + k_1)(\sigma_{\hat{y}}^2 + \sigma_y^2 + k_2)} \quad (7)$$

where μ represents the window mean intensity computed using a Gaussian filter, σ represents the window standard deviation, $\sigma_{\hat{y}y}$ is the window covariance between the predicted and real images, and $k_1 = 0.01$, $k_2 = 0.03$ are small constants to prevent division by zero. The SSIM is computed independently for each spectral channel, and the final imagewise SSIM score is obtained by averaging across all pixels and channels. The final $\text{SSIM}_{\text{score}}$ ranges from -1 to 1, where 1 indicates perfect structural similarity, 0 is not similar, and -1 is perfect dissimilarity.

As SSIM measures similarity, it is inverted and normalized to create the SDIM. SHAZAM computes two SDIM-based metrics: an imagewise anomaly score to flag if an image contains a potential hazard or not, and a pixelwise anomaly map to map the hazard. The imagewise anomaly score is calculated by inverting and normalizing the global SSIM score into a 0 to 1 range as follows:

$$\text{SDIM}_{\text{score}} = \frac{1 - \text{SSIM}_{\text{score}}}{2} \quad (8)$$

where $\text{SSIM}_{\text{score}}$ is the average SSIM across all pixels and channels.

A score of 0 means that the predicted image is identical to the true image. Scores increasing from 0 to 0.5 indicate increased differences between the images and, therefore, a higher indication of an anomaly. Scores between 0.5 and 1 indicate that the images are negatively correlated and dissimilar.

For the anomaly map, SHAZAM uses a few more steps. The anomaly map first computes the SSIM of each pixel across each spectral channel using the aforementioned local-window approach. It then calculates the average of these channelwise

maps to give a single score per pixel, and inverts them. The resulting values are clamped between 0 and 1, before being squared to suppress the background and emphasize the major structural differences as follows:

$$\text{SSIM}_{\text{avg}} = \frac{1}{C} \sum_{c=1}^C \text{SSIM}_c(\hat{y}, y) \quad (9)$$

$$\text{SDIM}_{\text{map}} = \max(0, \min(1, 1 - \text{SSIM}_{\text{avg}}))^2 \quad (10)$$

where C represents the number of channels. Overall, this two-pronged approach provides both a quantitative measure of unexpected change in the ROI, and a distinct map that highlights where these changes have occurred.

F. Seasonal Threshold

Instead of a flat threshold for anomaly detection, a seasonal threshold is used to account for variations in SIU-Net's performance throughout the year. This is justified by the difficulty of reconstructing seasons with significant structural differences from the baseline image. For example, predicting the first snowfall in winter is challenging, as snow accumulation patterns vary annually and obscure landmarks differently each year. This makes it harder for SIU-Net to model the ROI, increasing the average anomaly score during this period. Therefore, a cyclical seasonal threshold is proposed that changes across the year.

To calculate the seasonal threshold, a linear regression is first fitted to the anomaly scores from the training period using circular day-of-year coordinates to ensure continuity across the year boundary as follows:

$$\hat{m}_s(t) = a_1 \cdot t_{\sin} + b_1 \cdot t_{\cos} + c_1 \quad (11)$$

where $\hat{m}_s(t)$ represents the expected anomaly score for day t .

A second linear regression is fitted to the monthly standard deviations of the anomaly scores, capturing the seasonal variation in SIU-Net's performance as follows:

$$\hat{s}_s(t) = a_2 \cdot t_{\sin} + b_2 \cdot t_{\cos} + c_2 \quad (12)$$

where $\hat{s}_s(t)$ represents the expected standard deviation for day t .

The seasonal threshold curve is then defined as follows:

$$\tau(t) = \hat{m}_s(t) + 1.64 \times \hat{s}_s(t) \quad (13)$$

where $\tau(t)$ is the threshold value for day t . The multiplier of 1.64 corresponds to the 95th percentile of a normal distribution (one-tailed), allowing the threshold to capture approximately 95% of the training dataset's normal variations. Scores that exceed this threshold are classified as anomalies. Fig. 6 illustrates this adaptive threshold applied to the SNP dataset.

IV. EXPERIMENTAL SETUP

A. Datasets

Four datasets were collected through the Sentinel-2 satellite constellation and extracted using Sentinel-Hub¹ and its associ-




Fig. 6. Example of the seasonal threshold as calculated on training images from the SNP dataset covered in the following section.

ated Python API. These datasets cover different geographical regions of interest with various types of hazards. The images all use Level-2a processing, which means that the pixel values have been radiometrically corrected to provide bottom-of-atmosphere surface reflectance. From the available 13 spectral bands, only the 10 bands with 10 m or 20 m resolution are used, excluding the three 60 m resolution bands, which are primarily used for atmospheric correction. The SNP, Gran Chaco, and Lake Copeton datasets are further downsampled to 30 m pixel resolution, as this is sufficient for detecting their respective hazards while allowing coverage of larger regions of interest. The geographical locations of all datasets and hazard examples are shown in Fig. 7. Further details are provided in the following sections, with their key attributes shown in Table I.

1) *Sequoia National Park (SNP)*: This dataset was collected over a forested mountain region in the Sequoia National Park in California, the United States of America. This is a unique area that contains both wildfires and extreme out-of-season snowfall throughout the year. Detecting out-of-season snowfall is useful for nearby communities, which are at risk of flooding when the snow melts. Images captured when snow depth was above or below the maximum and minimum of the training period were labeled anomalies, based on data obtained via the National Water and Climate Center for the Leavitt Meadows ground station near to the ROI [95]. The wildfires were identified by visual inspection of RGB and infrared images. Images with an average NBR below the training minimum were also labeled as anomalies, as they indicate large amounts of destroyed vegetation in the ROI (i.e., burned regions).

2) *Gran Chaco*: Gran Chaco is the second largest forest in South America, stretching into Argentina, Bolivia, Paraguay, and Brazil. It is a hot and dry forest that is sparsely populated and is under constant threat of deforestation to create agricultural land [96], [97]. Remote detection of deforestation allows local authorities to intervene and reduce further forest degradation. Deforestation is the primary anomaly present, with some minor fires used for crop burning. The test images were labeled as anomalies by visual inspection in combination with outliers identified using the NDVI. This dataset has a consistent increase in deforestation over time and is therefore highly unbalanced, as almost all test images are considered anomalous.

¹[Online]. Available: <https://www.sentinel-hub.com/>



Fig. 7. Geographical locations of the four datasets/ROIs. Examples of their respective hazards are shown.

TABLE I
DATASET ATTRIBUTES FOR EACH ROI

Dataset Attributes	Sequoia National Park (SNP)	Gran Chaco	Lake Copeton	Lismore
Spatial Res. (Pixels)	1024x1024	1024x1024	512x512	1024x1024
Pixel Size	30m	30m	30m	10m
Hazards	Wildfires, Extreme Snowfall	Deforestation	Algal Blooms, Drought	Floods
Training Period	Jan 2017 to July 2020	Jan 2018 to June 2020	Jan 2022 to Apr 2024	June 2017 to June 2020
No. of Training Images	186	56	78	83
Test Period	Aug 2020 to Aug 2023	July 2020 to Apr 2024	Jan 2020 to Dec 2021	July 2020 to Aug 2024
No. of Test Images	250	110	61	105
Test Images with Hazards	187	108	56	24
Hazard Image Percentage	74.8%	98.2%	91.8%	22.9%

3) *Lake Copeton*: Lake Copeton is a lake bounded by Copeton dam in regional New South Wales, Australia. Lake Copeton provides freshwater to nearby communities for agricultural, household, and recreational purposes and is also home to aquatic life [98]. The use of upstream fertilizers has led to algal blooms and reduced water quality, especially during heavy rainfall. These water quality drops pose health risks to the ecosystem and community if not properly managed. In addition, the region is prone to drought and significant variations in the amount of water stored in the lake.

Images with algal blooms and reduced water quality were partially labeled with periods corresponding to water quality alerts from the government organization that manages Copeton Dam, WaterNSW [99]. Visual inspection of RGB and normalized difference chlorophyll index [100] was also used to label images with visible water disturbances, such as large sediment clouds. Images were labeled as drought when Copeton Dam had less than 50% water capacity [101].

4) *Lismore*: Lismore, a town in regional New South Wales, Australia, is prone to flooding, having experienced more than ten flood events in the last two decades [102]. Although satellite images are not necessarily the most effective method for flood detection in populated areas, they provide valuable spatial information for mapping flood impacts. The 2022 floods altered the landscape by causing extensive damage to infrastructure and agricultural land, with an estimated cost of approximately \$1 billion AUD [103]. While active flooding is captured in only a few images due to its brief duration, postflood effects are visible through damaged structures and altered land surfaces. These floods are visually confirmed and documented in Lismore Council's records.

The images are labeled as hazardous from the February/March 2022 floods until January 2023, as visual analysis of the satellite imagery showed larger flood-induced landscape alterations diminishing around this time. The dataset serves as an effective test case for SHAZAM's ability to detect both immediate flood

impacts and persistent landscape changes. The area also undergoes gradual changes in land cover due to urban and agricultural development, adding complexity to the dataset as these changes may be detected as potential anomalies.

B. Implementation and Comparative Models

1) *SHAZAM*: SHAZAM, using the SIU-Net architecture, was trained for 20 epochs with a batch size of 32. For all datasets, the training images were cut into 32×32 patches, with 90% used to train SIU-Net for patch translation, and 10% randomly selected to validate the model's translation performance. The L1 loss, also known as the MAE, was used as the loss function. An initial learning rate of $1e-4$ with the Adam optimizer was used (with default parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$). The learning rate was progressively reduced by a factor of 0.1, whenever the validation loss plateaued for 3 epochs. The minimum learning rate was set to $1e-7$. Once trained patchwise, SHAZAM was implemented imagewise to the training period so that the seasonal threshold could be estimated using the training image anomaly scores. Finally, SHAZAM was applied to the test period for final evaluation (as shown in the inference stage of Fig. 2). All models, including SHAZAM, were implemented in Python (with PyTorch Lightning² used for deep learning). The layer weights are initialized with the PyTorch defaults.

2) *cVAE*: A conditional variational autoencoder (cVAE) serves as a baseline model, chosen for its ability to learn complex data distributions based on conditional inputs [104], [105]. The cVAE employs a symmetric encoder-decoder structure with residual connections, where the encoder progressively down-samples the input through alternating convolutional and residual blocks while increasing channel depth (32, 64, 128). The output is then mapped to a normalized 256-D latent space using the standard VAE reparameterization trick.

The latent space is conditioned on four variables: the cyclical day-of-year encodings (sin, cos) and the scaled patch positions (row, column). These variables directly enable the model to capture seasonal patterns and regional variations across the ROI. The decoder mirrors the encoder's architecture, but uses transpose convolutions to upsample the latent features to the original patch size. During inference, only the decoder is used to generate patches given temporal-spatial conditions. Training is guided by an L1 reconstruction loss of the input and the Kullback–Leibler divergence ($\beta = 0.1$) to balance generation quality with a Gaussian latent space. For training, a batch size of 32 was used alongside 20 epochs. cVAE uses the same data preprocessing (including normalization) and optimization strategy as SHAZAM. The training phase begins with an initial learning rate of $1e-4$, decreasing by 0.1 when the validation loss plateaus for 3 epochs. For inference, the average L1 loss per image is used for ROI hazard detection, and the L1 loss per pixel is used for mapping. The hazard detection threshold is the mean L1 loss of the training images plus 1.64 standard deviations, following the theoretical justification as SHAZAM's threshold.

3) *RaVAEn*: RaVAEn provides the closest fit to the objective, given its direct design for change monitoring to detect natural hazards/disasters in Sentinel-2 time-series patches [55]. RaVAEn uses a VAE to compress patches into a 128-D latent space, and then uses the minimum cosine distance between the current patch's latent features and historical patch latent features to score events. The three latest patches are used to detect change events.

Given that RaVAEn was designed for 32×32 pixel patches and the ROI images are much larger than this, two variants are implemented. The first, RaVAEn-Local, calculates the minimum cosine distance for each patch as in the paper (comparing the current patch to the three historical patches). The minimum cosine distance is then used to represent all pixels in the 32×32 patch, which is repeated for all patches in the ROI to create the anomaly heatmap. The average score across the anomaly heatmap is used to flag a hazard in the region. The second variant, RaVAEn-Global, computes and stores the cosine distances for the current patch and each of the three preceding patches. These values are used to create three anomaly heatmaps for each preceding image of the ROI. The heatmap with the lowest average score is then used to detect and map anomalies. The hazard detection threshold is the mean L1 loss of the training images plus 1.64 standard deviations, following the theoretical justification as SHAZAM's threshold. Cubic convolution is used to upsample the heatmaps to match the image resolution for visual inspection. For training, a batch size of 32 was used alongside 10 epochs, with a fixed learning rate of $1e-3$ for the Adam optimizer.

The s2cloudless cloud filtering algorithm was not used due to the low cloud coverage in the datasets, and because the algorithm risks filtering out snow-based pixels. This is particularly important given that the extreme and out-of-season snowfall is a key hazard in the SNP dataset. Nonetheless, the handling of cloud coverage remains a crucial requirement for optical SITS monitoring systems and is discussed in further detail in Section V-E.

4) *COLD*: The COLD algorithm provides a good point of comparison with disturbance detection methods, as it is designed to detect various types of disturbance using multiple bands [14]. Designed for Landsat SITS, COLD models each band's time series independently for each pixel in the image. It captures trends and seasonality using a limited Fourier series. The author's implementation in the PYCOLD library is used.³

Several adaptations were made to implement COLD with Sentinel-2 data for immediate hazard detection and mapping. First, only bands that closely match Landsat wavelengths were used (blue, green, red, near-infrared, short-wave infrared 1, and short-wave infrared 2, with no thermal band being used). Second, to balance computational efficiency with accuracy, the analysis was performed on the mean values of 8×8 pixel patches rather than individual pixels. Disturbance detection methods usually need HPCs to process an ROI promptly, which are unlikely to be available for constant monitoring. Third, instead of using

²[Online]. Available: <https://github.com/Lightning-AI/pytorch-lightning>

³[Online]. Available: <https://github.com/GERSL/pycold>

COLD's approach of multiple observations to confirm a disturbance, the mean disturbance probability is used as an anomaly score. This allows COLD to be tested for detecting and mapping hazards on the first observation. Finally, like RaVAEn, COLD's cloud filtering is omitted because of the low cloud coverage and the need to detect snow-based hazards.

When evaluating each test image, COLD uses all available historical images from the datasets to model the time series of each 8×8 patch. This includes both the complete training dataset and any previous test images. However, COLD could not be evaluated on the Lake Copeton dataset, which lacked sufficient data prior to the test period. Since all SITS training data are required for modeling, it is not feasible to calculate a threshold similar to the approaches used by SHAZAM, cVAE, and the RaVAEn variants. Therefore, COLD is compared using only thresholdless metrics (discussed in the following section). Like RaVAEn, cubic convolutions are used to upsample the heatmaps.

C. Metrics

The precision, recall, and F1 scores are used to evaluate performance of each model when using a specific threshold τ . Precision is the ratio of true positives to all identified hazards, which assesses the accuracy of the model for hazard detection and the risk of false alarms. Minimizing false alarms (and maximizing precision) is crucial to avoid wasting resources on unnecessary hazard confirmations or responses. Recall measures correctly identified hazards against the total number of true hazardous images, considering false negatives (or missed hazards). Recall is crucial for hazard detection, as missing hazards can have severe consequences. The F1 score combines both precision and recall into a single aggregated metric as follows:

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}} \quad (14)$$

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} \quad (15)$$

$$\text{F1 Score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (16)$$

The F1 score can be considered a more accurate measure than accuracy, as it accounts for the imbalance between anomalous and normal images in the test dataset. In addition, the area under the precision-recall curve (AUPRC) is used to evaluate the models across a range of possible thresholds. This is a robust secondary metric that assesses general detection performance without a specific threshold value, and is also unaffected by the imbalance between change events. Precision, recall, F1 score, and AUPRC range from 0 to 1, with higher values denoting better performance.

An important note is that SHAZAM's AUPRC is calculated using the residuals after removing the seasonal threshold. Since the seasonal threshold models SHAZAM's error rates throughout the year (mean fit plus 1.64 times the monthly standard deviation fit), removing it leaves the residual scores that determine if an image is anomalous (greater than 0) or normal (less than zero). This allows the AUPRC to evaluate performance across

a range of constant thresholds, providing a fairer assessment inline with the comparative models while allowing SHAZAM to benefit from its complete seasonal modeling.

As many of the datasets are heavily imbalanced, a "coin-toss" random model is used as a baseline against which learning-based models perform the task. This random model simply guesses with a 50% chance if any given image contains a hazard or not. The metrics for this random model can be theoretically calculated, yielding its precision, recall, and F1 score as follows:

$$\text{Precision}_{\text{rand}} = \frac{p \cdot P}{p \cdot P + p \cdot (1 - P)} = P \quad (17)$$

$$\text{Recall}_{\text{rand}} = \frac{p \cdot P}{p \cdot P + (1 - p) \cdot P} = 0.5 \quad (18)$$

$$\text{F1 Score}_{\text{rand}} = \frac{2 \cdot P \cdot p}{P + p} = \frac{P}{P + 0.5} \quad (19)$$

where $P \in [0, 1]$ is the fraction of positive samples in the dataset (i.e., the fraction of images labeled as hazards), and $p = 0.5$ is the probability of the random model classifying an image as having a hazard. Given that the precision of the random model is constant at P regardless of the classification threshold, and thus remains fixed for all values of recall, the AUPRC is simply the same as the precision as follows:

$$\text{AUPRC}_{\text{rand}} = P. \quad (20)$$

D. Experiments and Hardware

Change monitoring differs from change detection by aiming to flag events at the image level, rather than mapping individual changed pixels. Given this focus on event detection for an ROI, the first experiment evaluates SHAZAM and the comparative models in hazard detection at the image level across all datasets. The detection performance of each model is evaluated using the quantitative metrics covered in the previous section, and a qualitative analysis of the anomaly heatmaps is used to evaluate the mapping ability. A time-series analysis of the anomaly scores is also conducted on the SNP dataset (the most complex dataset) to better understand the edge cases where SHAZAM still requires improvement, highlighting missed hazards (false negatives) and false alarms (false positives).

The second core experiment is an ablation study that evaluates SHAZAM's core components by systematically removing positional encodings, cyclical seasonal encodings, the structural difference module, and the seasonal threshold. Seasonal encodings are removed by replacing SIU-Net and the seasonal threshold's sine and cosine day-of-year embeddings with the normalized day-of-year (0 for January 1st to 1 for December 31st). To ablate the structural difference, it is replaced by the MAE. This analysis helps validate design choices and explores potential simplifications of the architecture.

All experiments are conducted onboard a Windows desktop with an AMD Ryzen 9 3900X 12-Core Processor, 64 GB of RAM, and an NVIDIA GeForce RTX 3080 graphics card with 10 GB of VRAM. This is with the exception of COLD, which was implemented in a Google Colab environment (Linux), with access to 51 GB of RAM and an 8-core CPU.

TABLE II
DETECTION RESULTS

Model	#Params	Sequoia National Park (SNP)				#Params	Gran Chaco			
		F1	Precision	Recall	AUPRC		F1	Precision	Recall	AUPRC
Random Guess	-	0.599	0.748	0.5	0.748	-	0.662	0.982	0.5	0.982
cVAE	2.6M	0.245	0.470	0.169	0.627	2.6M	0.230	1.000	0.130	0.999
COLD	41.9M	-	-	-	0.719	41.9M	-	-	-	0.995
RaVAEn-Local	617K	0.182	0.864	0.102	0.838	617K	0.108	1.000	0.057	0.967
RaVAEn-Global	617K	0.181	0.826	0.102	0.821	617K	0.108	1.000	0.057	0.967
SHAZAM	473K	0.771	0.767	0.775	0.762	473K	0.728	0.969	0.583	0.987

Model	#Params	Lake Copeton				#Params	Lismore			
		F1	Precision	Recall	AUPRC		F1	Precision	Recall	AUPRC
Random Guess	-	0.647	0.918	0.5	0.918	-	0.314	0.229	0.5	0.229
cVAE	2.6M	0.324	0.917	0.196	0.918	2.6M	0.069	0.200	0.042	0.204
COLD	-	-	-	-	-	41.9M	-	-	-	0.391
RaVAEn-Local	617K	0.486	1.000	0.321	0.949	617K	0.061	0.111	0.042	0.226
RaVAEn-Global	617K	0.172	1.000	0.094	0.935	617K	0.054	0.077	0.042	0.246
SHAZAM	473K	0.881	0.906	0.857	0.928	473K	0.466	0.304	1.000	0.389

Bold highlights the best results.

V. EXPERIMENTAL RESULTS AND ANALYSIS

A. Hazard Detection Comparison

The detection results of SHAZAM and the comparative models are shown in Table II. Since the Lake Copeton dataset’s test period predates the training period, COLD was not used because it needs sequential historical data for change detection.

Table II demonstrates SHAZAM’s superior performance on the SNP dataset, achieving both the highest F1 score and recall. Although its precision is comparable to the random baseline, SHAZAM’s substantially higher recall than all the other models indicates an enhanced ability to detect hazards events with minimal misses. In contrast, the RaVAEn variants achieve the highest precision and AUPRC, showing excellent performance when minimizing false alarms, but their very low recall indicates they miss most hazardous images. The disconnect between RaVAEn’s strong AUPRC and poor recall and F1 score suggests the potential for improvement through threshold optimization. However, SHAZAM’s overall stronger and balanced metrics demonstrate that the use of its seasonal threshold provides an effective empirical tradeoff between false alarms and missed hazards.

The results of the Gran Chaco dataset follow a similar pattern, with SHAZAM demonstrating superior performance through a higher F1 score and recall. Although SHAZAM has marginally lower precision compared to other models, the precision is notably high across all approaches. This aligns with the region’s characteristics of persistent change from gradual deforestation, making false alarms almost impossible. The extremely low recall scores for most models indicate that they miss many hazards due to the default threshold (mean + 1.64 standard deviations) being overly conservative. While the high AUPRC scores suggest strong hazard separation capabilities across all models, this is less meaningful given that setting a low threshold in this region naturally yields high performance due to the high hazard percentage.

SHAZAM demonstrates superior performance across all metrics for the Lake Copeton dataset, except precision. Its F1 score

exceeds the random baseline by 0.234 and exceeds the next best model, RaVAEn-Local, by 0.395, marking a substantial improvement. There is a notable discrepancy between the recall and the F1 scores of the two RaVAEn variants. Despite similar AUPRCs and perfect precision values, investigation revealed that RaVAEn-Global’s threshold was significantly higher due to elevated anomaly score means and standard deviations in the training set. Further analysis showed this was caused by the Lake Copeton dataset’s more prominent cloud coverage within images, and also extended gaps between images due to a more irregular sequence. RaVAEn-Global’s higher sensitivity to these factors results in more pronounced seasonal changes, leading to inflated anomaly scores during threshold calculation. This finding emphasizes the importance of managing cloud coverage, seasonality, and the irregularity of optical SITS.

SHAZAM demonstrates superior performance across all metrics for the Lismore dataset, with COLD being the only model to achieve comparable AUPRC. SHAZAM’s perfect recall indicates that it correctly detects all instances of hazardous images, though its lower precision suggests a tendency toward false alarms. While both Lismore and Gran Chaco experience persistent changes, Lismore’s changes are predominantly non-hazardous, stemming from routine human activities such as constructing more buildings and agricultural modifications. This makes false alarms more likely as these regular land cover changes are flagged despite not being hazardous. For regions experiencing such persistent nonhazardous changes, updating the baseline image or using another method to capture trend-based changes could be a viable solution. The comparative models perform significantly worse than the random baseline in terms of F1 score, precision, and recall. While threshold optimization could marginally improve their results, SHAZAM clearly demonstrates more robust performance.

Across all datasets, SHAZAM consistently demonstrates superior detection performance compared to the cVAE and RaVAEn variants, with F1 score improvements ranging from 0.066 on the Gran Chaco dataset, to 0.234 on the Lake Copeton dataset. While SHAZAM achieves comparable or superior precision scores, its stronger performance is primarily driven by

higher recall scores, indicating greater effectiveness in detecting hazards while reducing missed detections. COLD could only be evaluated using AUPRC due to dataset limitations preventing threshold calculation, but within the AUPRC scores, SHAZAM typically matched or slightly trailed other models (except Lake Copeton where it leads). This consistent performance, coupled with the fact that cVAE and RaVAEn variants perform below the random baseline, highlights both the efficacy of the seasonal threshold and the broader importance of threshold selection in change monitoring. The fundamental limitations of these competing approaches are explored in detail through visual comparison in the following section. SHAZAM's lightweight architecture makes it resource and cost-efficient, an essential requirement for practical monitoring systems. Overall, these results demonstrate that SHAZAM successfully balances the critical tradeoff between false alarms and missed hazards while maintaining robust detection performance across diverse geographical regions and hazard types.

B. Visual Mapping Evaluation

The mapping results are grouped by dataset. Fig. 8(a) compares the heatmaps of each model on the SNP dataset, using a live wildfire and out-of-season snowfall as examples. Fig. 8(b) shows a small section of deforestation mapping on the Gran Chaco dataset, while Fig. 8(c) shows an algal bloom and drought in Lake Copeton. Finally, Fig. 8(d) compares the flood mapping capabilities of each model.

SHAZAM has distinctively strong performance when mapping both the wildfires and the out-of-season snowfall in Fig. 8(a). For wildfire detection, SHAZAM achieves pixel-precise mapping of both burned areas and active fires/smoke plumes, effectively differentiating between active fire locations and their effects. While COLD's disturbance probability map identifies affected areas, it lacks the ability to distinguish active fire fronts. The cVAE output, though highlighting active fires, fails to maintain the scene's spatial structure, making it difficult to differentiate fire-related features from the background. RaVAEn-Global's strength in capturing dynamic events is evident in its rough detection of active fires and smoke, but its lower resolution output makes it much less precise, and its dynamic design misses crucial burned regions. RaVAEn-Local, designed for patchwise scoring, introduces considerable noise by incorporating signals from all three prior temporal images, in contrast to RaVAEn-Global's more focused single-image comparison approach.

For the out-of-season snowfall, the conventional detection methods struggle. The COLD algorithm does not identify the snow coverage, while cVAE, despite leveraging the high reflectance of snow, is unable to suppress the background terrain. The RaVAEn variants demonstrate some success in delineating the snow-front boundaries, but struggle to map this out-of-season anomaly. In contrast, SHAZAM is very effective at suppressing the background while accentuating the unexpected snow, demonstrating a strong ability to model seasonality.

By analyzing deforestation mapping in the Gran Chaco dataset [see Fig. 8(b)], conventional methods show varying

limitations. COLD sporadically maps some of the new deforestation within the cyan box while effectively suppressing the background, but fails to reveal true structure. cVAE highlights all deforested areas, including those from the training dataset, due to its inability to generate accurate structural representations of the ROI. Both RaVAEn variants achieve good background suppression, but are unable to map the newly deforested areas. RaVAEn-Global's mapping is slightly more visible than RaVAEn-Local but the deforestation patterns are still not visible, a limitation of its coarse patchwise resolution.

SHAZAM demonstrates superior detection of newly deforested regions, with clear mapping of recent land-use changes. Notably, it identifies anomalies within existing agricultural fields that correspond to recent reclearing or burning activities. The detection of two fields outside and to the bottom right of the cyan box, which were cleared before the test period, reveals an important characteristic of SHAZAM's methodology. As the baseline represents an average image across the entire training period, SHAZAM's ability to detect structural changes depends on this reference, suggesting that some recent features may not be captured if they are partially incorporated into the baseline.

The Lake Copeton dataset demonstrates SHAZAM's ability to map algal blooms, a type of hazard that the other models had not previously evaluated [see Fig. 8(c)]. cVAE maps the blooms within the water but struggles to distinguish them from the background. The RaVAEn variants show some ability to map the blooms, but at the noticeably lower resolution. Although they capture larger clusters of algae, these methods miss the finer strains present in the scene. SHAZAM, in contrast, precisely maps all algal blooms present on the surface. SHAZAM maintains a high spatial resolution that captures both the larger clusters, and the finer scale threads between them and along the lake's edge.

The drought example shows marked differences in mapping and detection capabilities (see Fig. 8(c), bottom row). cVAE and both RaVAEn variants fail to map the drought conditions. While cVAE's failure stems from poor ROI structure representation, RaVAEn's limitation is inherent to its design focus on rapid "change events." By comparing the current satellite image to only a small window of recent images, RaVAEn lacks the historical context to establish normal water levels for the lake, making it unsuitable for detecting and mapping hazards that gradually develop, such as drought. SHAZAM, using its baseline-derived ROI with seasonal translation, clearly delineates the drought-affected region. The increased background signal in the anomaly heatmap is consistent with drought conditions, reflecting the broader impact of extreme dry seasons on the surrounding terrain.

The Lismore dataset demonstrates the capability of each method in flood mapping [see Fig. 8(d)]. COLD, despite not being specifically designed for flood detection or immediate hazard response, performs notably well in identifying flooded rivers and shows clusters of anomalous pixels in inundated areas, though it fails to capture the complete extent of the flood. cVAE is ineffective on this dataset, while the RaVAEn variants identify the larger flooded regions but




Fig. 8. Change/hazard heatmap examples of different hazards from each dataset. (a) SNP dataset heatmap comparison. The top row shows heatmaps for a bushfire event, while the bottom row shows out-of-season snowfall. (b) Gran Chaco dataset heatmap comparison, showing deforestation mapping. The cyan box highlights the primary area of new deforestation since the training period. (c) Lake Copeton dataset heatmap comparison, showing flood mapping. COLD is absent from this dataset due to no pretest images in the Lake Copeton dataset. (d) Lismore dataset heatmap comparison, showing flood mapping.




Fig. 9. Time-series plot of SHAZAM-calculated anomaly score residuals for all images in the SNP dataset, obtained by subtracting the seasonal threshold from each image's anomaly score. The numbers indicate selected examples of outliers in the training set and false positives/negatives in the test set, and corresponding images are provided. The training outliers and false positives are attention-weighted by SHAZAM's calculated anomaly heatmap, while the false negatives simply show a true satellite image. (a) Train anomaly: Wildfire. (b) Train anomaly: Snow. (c) False positive: Snow. (d) False negative: Burned region. (e) False negative: Burned region. (f) False positive: Burned region.

miss finer details due to their characteristic lower resolution. RaVAEn-Local shows a better ability to suppress the background in this case. SHAZAM achieves comprehensive flood mapping with high spatial resolution, clearly mapping all flood-affected areas.

Across diverse datasets and hazard types, SHAZAM demonstrates consistently superior mapping capabilities compared to existing methods. Although COLD shows some ability to map floods, fires, and some deforestation, its reliance on multiple disturbance observations often results in partial or missed hazard detection. cVAE exhibits generally poor performance across all hazards; despite some understanding of seasonality, it struggles to learn the spatial structure of the scene and suppress background features. The RaVAEn variants, though effective at background suppression and mapping abrupt changes, are limited by their coarse resolution and reliance on recent temporal context, particularly evident in their inability to detect gradual changes such as drought and deforestation. Although useful for detecting dynamic change events, their lack of seasonal understanding is further demonstrated by their inability to detect out-of-season snowfall.

SHAZAM effectively combines seasonal modeling with high spatial resolution to accurately map all example hazards. This includes immediate events such as floods and fires, to more gradual changes such as drought and deforestation. This consistent performance across diverse scenarios, coupled with its ability to capture both fine-scale details and broader patterns, demonstrates SHAZAM's versatile and seasonally-aware approach to hazard mapping and detection.

C. Time-Series Analysis and Sample Edge Cases

The time-series plot in Fig. 9 reveals key findings about SHAZAM from both the training and test data in the SNP dataset. Despite the assumption that the training data is normal, anomalies are still flagged, which includes potential hazards such as the small wildfires seen in (a). This supports the notion that training datasets need not be hazard-free to be detected, as they are assumed to be rare anomalies and unexpected events. Creating perfect, hazard-free datasets is impractical, especially with the growing abundance of satellite images and the desire to avoid manual labeling by leveraging self-supervised learning. This also demonstrates SHAZAM's power for retrospective event discovery even when no manual filtering has been applied. Another anomaly is snow in (b), often flagged due to its high reflectivity and variable distribution across the scene, especially when early and late-season snowfall occurs unexpectedly. Expanding the dataset using decades of historical images, such as Landsat or harmonized Landsat-Sentinel, is essential to further enhance seasonal modeling.

The test period shows seasonal clusters of false negatives and positives, particularly around the transition to spring, also due to unpredictable or longer lasting snowfall. An example is (c), where SHAZAM focuses on snow being anomalous even though it is winter and within the expected snow-depth range. However, the snow depth recording station is only a proxy at a different location, potentially misidentifying anomalies. The false negatives in (d) and (e) include subtle hazards such as burned regions after wildfires. An intriguing case is (f), where

TABLE III
SHAZAM ABLATION STUDY—F1 AND AUPRC SCORES

Model	SNP		Gran Chaco		Lake Copeton		Lismore	
	F1	AUPRC	F1	AUPRC	F1	AUPRC	F1	AUPRC
No Position	0.762	0.788	0.714	0.993	0.891	0.923	0.460	0.390
Linear Time	0.490	0.684	0.579	0.995	0.917	0.954	0.495	0.392
MAE Score	0.705	0.755	0.876	0.993	0.870	0.965	0.512	0.414
Flat Threshold	0.378	0.666	0.468	0.987	0.865	0.958	0.533	0.355
SHAZAM	0.771	0.762	0.728	0.989	0.881	0.928	0.466	0.389

Bold highlights the best results.




Fig. 10. Comparison of SHAZAM's generated output with and without positional encodings. The cyan box shows an area where removing position encodings leads to a coarser representation of expected snow cover.

burned areas are clearly flagged as false positives, possibly due to longer term scarring from the 2022 wildfires.

Overall, SHAZAM has strong monitoring performance when detecting hazards and detecting unlabeled training anomalies. However, this time-series analysis highlights clear opportunities to further improve seasonal modeling and sharpen the separation between genuine anomalies and normal variations. It also underscores the challenge of hazard labeling: anomalies may stem from normal seasonal changes or long-lasting posthazard signatures. Integrating more thorough in-situ observations from an ROI with additional expert validation would strengthen the ground truth and further sharpen the distinction between true anomalies and normal variability.

D. Ablation Study

The effects of removing positional encodings, seasonal encodings, SSIM score, and the seasonal threshold are shown in Table III. Removing the position encodings shows minimal degradation in F1 scores and even marginal improvements in AUPRC (a 0.026 increase for SNP), qualitative analysis reveals that positional encodings contribute to more precise spatial representations (see Fig. 10), also shown by slightly improved MAE and SSIM scores in the training phase. This suggests that the position encodings may improve feature representation at a local level, though their impact on final classification metrics is dataset-dependent. For smaller SITS datasets, there may be a risk of overfitting.

Replacing the cyclical day-of-year encodings with linear temporal representations reveals more noticeable impacts on model performance. While marginal improvements are seen in the Lake Copeton and Lismore datasets (F1 increases of 0.036 and 0.029, respectively), there is a significant performance drop in SNP (F1 decrease of 0.281 and AUPRC decrease of 0.078) and Gran

Chaco (F1 decrease of 0.135). Overall, these results show that cyclical day-of-year encodings provide a robust approach for incorporating seasonality across diverse locations and ROIs, and are particularly effective in regions with more pronounced seasonal changes (such as the SNP dataset).

The comparison between SSIM and MAE scoring mechanisms provides important insights into SHAZAM's detection framework. Using MAE for anomaly scoring achieves marginally lower AUPRC scores on SNP (decrease of 0.007) and Gran Chaco (decrease of 0.004), while showing improved performance on Lake Copeton (increase of 0.037) and Lismore (increase of 0.025). The F1 scores reveal a tradeoff: MAE demonstrates improved performance on Gran Chaco (increase of 0.148) and Lismore (increase of 0.046) but shows decreased effectiveness on SNP (decrease of 0.066) and Lake Copeton (decrease of 0.011). These results validate SHAZAM's fundamental approach of comparing generated and true images, while opening up further research into more optimized scoring methods. Nevertheless, the SSIM-based structural difference module provides distinct advantages through its ability to capture local structural differences while suppressing background variations when mapping hazards.

Using a flat threshold, which removes the seasonal threshold, further demonstrates the importance of seasonal modeling within SHAZAM. The substantial reduction in F1 scores on the SNP (decrease of 0.393) and Gran Chaco (decrease of 0.260) datasets, alongside generally reduced performance across all dataset F1 scores and AUPRC scores, validates the need for an adaptive threshold that accounts for seasonal variations. The exception is the Lismore dataset, which shows a modest improvement (increase of 0.067) in the F1 score, suggesting that simpler threshold strategies might be sufficient in some cases, although at the cost of reduced generalization. Collectively, these ablation results support SHAZAM's architectural design while providing deeper insights into the contribution of each component. While individual components may show advantages in specific contexts, SHAZAM's full architecture provides a theoretically grounded and empirically robust approach to self-supervised change monitoring for hazard detection and mapping.

E. Broader Limitations and Future Work

Although SHAZAM demonstrates strong performance for the detection and mapping of various types of hazards, several key challenges remain to create an ideal operational hazard monitoring system. A key limitation is the management of cloud

coverage, as the current evaluation uses datasets with less than 10% cloud cover. While existing methods such as COLD and RaVAEn implement pixelwise cloud masking and exclusion in their original implementations, this approach results in data loss and would greatly impair SHAZAM's structural difference module. Potential solutions include using SAR images, which can penetrate clouds and operate at night. Recent generative methods for converting SAR to multispectral images [40], [106] can enhance cloud-free multispectral SITS by increasing data availability and effective temporal resolution, while still leveraging the more detailed information available in optical images.

SHAZAM's intra-annual seasonal modeling also has limitations, as it does not capture gradual changes driven by longer term weather patterns (e.g., El Niño and La Niña). Future work could explore modeling these longer term patterns and trends, even including multimodal approaches that integrate additional data sources to supplement the optical imagery. The temporal frequency could also be increasing by using harmonized Landsat and Sentinel-2 data, although this challenge will naturally be reduced with more frequent image capture and the growth of remote sensing archives over time.

A third core challenge is scaling to multiregion or global coverage, an essential milestone for operational hazard monitoring. This requires models that generalize across different ROIs while maintaining a strong understanding of each region's spatial structure. Future work could explore injecting geographical coordinates alongside seasonality to model changes in land cover at diverse locations simultaneously. This video anomaly detectionesque approach could replace SIU-Net with a more scalable model, and the development of SITS foundation models could also enable generalized hazard detection while supporting other downstream tasks. However, computational efficiency should remain a critical consideration when designing scalable models, as it is a core requirement for operational deployment.

While SHAZAM effectively detects and maps hazards, its operational value would be significantly enhanced if it could also classify what the hazards are. Semisupervised and few-shot learning methods could leverage a small number of labeled images to develop hazard classification abilities, making efficient use of limited labeled data [107], [108]. Alternatively, zero-shot learning approaches could allow classification without labeled data from large, generalizable models [109], such as multimodal large language models (LLMs).

A critical direction for future work is onboard real-time monitoring for hazard detection and mapping. Although satellite-based systems provide global land coverage, they require significant time for data downlink and postprocessing. Performing onboard anomaly detection with less-processed data greatly reduces these delays. In addition, satellite multiday revisit times limit rapid response capabilities in situations requiring immediate mapping. Real-time hazard detection can leverage more agile platforms such as drones and aircraft that can be deployed rapidly. These platforms offer enhanced spatial resolution due to their proximity to the ground and can be readily equipped with advanced sensors, including hyperspectral cameras that provide rich spectral information for anomaly detection [110]. This combination of high spatial and spectral resolution, coupled with flexible deployment, enables

immediate and targeted hazard detection and mapping and opens new possibilities for real-time monitoring.

VI. CONCLUSION

The shift from change detection to change monitoring represents a crucial advancement in hazard detection and mapping. Existing methods often require hazard-specific labeled data, lack compatibility with irregular SITS observations, or fail to distinguish unexpected changes from normal seasonal variations. SHAZAM addresses these limitations through a self-supervised change monitoring approach that integrates seasonal modeling via SIU-Net, structural difference mapping and scoring, and a seasonally adaptive threshold. SHAZAM was evaluated alongside similar models on four datasets, which covered different ROIs with wildfires, extreme/out-of-season snowfall, deforestation, algal blooms, drought, and floods.

SHAZAM demonstrated superior hazard detection performance across all four datasets, with F1 score improvements between 0.066 and 0.234. This was primarily through its greater effectiveness at reducing missed hazards (higher recall). The consistently strong performance across all metrics demonstrated SHAZAM's ability to balance false alarms and missed detections in diverse geographical regions, while competing models performed below the random baseline. While other models could have benefited from dataset-specific threshold optimization, SHAZAM's theoretically grounded seasonal threshold required no optimization and was robust across all datasets. Notably, SHAZAM achieved this performance while being extremely lightweight, using the least number of learnable parameters among all comparative models (473 K).

SHAZAM demonstrated superior mapping capabilities through high spatial resolution, seasonal modeling, and the ability to suppress background features while accentuating changes caused by hazards. Although existing methods showed some success, they lacked spatial resolution, seasonal awareness, or missed mapping crucial hazard features. SHAZAM effectively mapped both immediate and gradual hazards, demonstrating versatility across various types of hazards and geographic locations.

A detailed residual-time-series study on the most challenging dataset (SNP) highlighted that SHAZAM still works when hazardous events are still present in "normal" training images. This aligns with the method learning to represent more common and expected seasonal behavior, and highlighting that self-supervised learning is a valid approach given that a perfectly hazard-free training dataset is not essential. Misclassifications are still concentrated around seasonal transitions or more subtle/smaller scale hazards, indicating that expanding the datasets over many more years or decades may further improve seasonal modeling. Edge-case inspection highlights that it is important not to assume that anomalies are always hazards; they can be benign extremes or lingering postevent signatures, underscoring the need for a greater depth of in-situ measurements and additional expert review to refine ground truth of hazard detection datasets.

The ablation study showed that altering individual components had both benefits and drawbacks, suggesting that further optimization is promising, but also that SHAZAM's current

architecture provides a theoretically grounded and empirically robust approach to hazard detection and mapping. Future work should address remaining challenges in cloud cover management, modeling long-term environmental changes, and scaling to multiregion or global coverage. The development of real-time detection capabilities and multimodal/LLM hazard classification offers promising directions that would enhance monitoring solutions for rapid insights into emerging hazards.

ACKNOWLEDGMENT


The authors would like to thank Dr. Lichao Mou for his advice, which was instrumental in the formation and early stages of this work. They would also like to thank Mr. Thomas Dujardin for his valuable technical advice and expertise in remote sensing and deep learning methods, and Dr. Pu Jin for his insights on ANDT.

REFERENCES

- [1] S. K. Abid et al., "Toward an integrated disaster management approach: How artificial intelligence can boost disaster management," *Sustainability*, vol. 13, no. 22, 2021, Art. no. 12560.
- [2] S. I. Seneviratne et al., "Weather and climate extreme events in a changing climate," in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, V. Masson-Delmotte et al., Eds. Cambridge, U.K. and New York, NY, USA, Cambridge Univ. Press, ch. 11, 2021, pp. 1517–1766, doi: [10.1017/9781009157896.013](https://doi.org/10.1017/9781009157896.013).
- [3] World Meteorological Organization, "Early Warnings For All - The UN Global Early Warning Initiative for the Implementation of Climate Adaptation: Executive Action Plan 2023-2027," Geneva, Switzerland, 2022. [Online]. Available: <https://library.wmo.int/idurl/4/58209>
- [4] S. N. Goward et al., "Semi-centennial of landsat observations & pending landsat 9 launch," *Photogrammetric Eng. Remote Sens.*, vol. 87, no. 8, pp. 533–539, 2021.
- [5] European Space Agency, "Missions-sentinel online," ESA, Paris, France, 2014. [Online]. Available: <https://sentinels.copernicus.eu/missions>
- [6] C. Rambour, N. Audebert, E. Koeniguer, B. L. Saux, M. Crucianu, and M. Datcu, "Flood detection in time series of optical and SAR images," *Int. Arch. Photogrammetry, Remote Sens. Spatial Inf. Sci.*, vol. 43, no. B2, pp. 1343–1346, 2020.
- [7] N. Casagli, E. Intrieri, V. Tofani, G. Gigli, and F. Raspini, "Landslide detection, monitoring and prediction with remote-sensing techniques," *Nature Rev. Earth Environ.*, vol. 4, no. 1, pp. 51–64, 2023.
- [8] P. P. D. Bem, O. A. de Carvalho Junior, R. F. Guimarães, and R. A. T. Gomes, "Change detection of deforestation in the Brazilian amazon using landsat data and convolutional neural networks," *Remote Sens.*, vol. 12, no. 6, 2020, Art. no. 901.
- [9] D. Rashkovetsky, F. Mauracher, M. Langer, and M. Schmitt, "Wildfire detection from multisensor satellite imagery using deep semantic segmentation," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 7001–7016, 2021.
- [10] I. Caballero, R. Fernández, O. M. Escalante, L. Mamán, and G. Navarro, "New capabilities of Sentinel-2A/B satellites combined with *in situ* data for monitoring small harmful algal blooms in complex coastal waters," *Sci. Rep.*, vol. 10, no. 1, 2020, Art. no. 8743.
- [11] A. Singh, "Review article digital change detection techniques using remotely-sensed data," *Int. J. Remote Sens.*, vol. 10, no. 6, pp. 989–1003, 1989.
- [12] M. Chi, A. Plaza, J. A. Benediktsson, Z. Sun, J. Shen, and Y. Zhu, "Big data for remote sensing: Challenges and opportunities," *Proc. IEEE*, vol. 104, no. 11, pp. 2207–2219, Nov. 2016.
- [13] C. Tao, J. Qi, M. Guo, Q. Zhu, and H. Li, "Self-supervised remote sensing feature learning: Learning paradigms, challenges, and future works," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5610426.
- [14] Z. Zhu et al., "Continuous monitoring of land disturbance based on landsat time series," *Remote Sens. Environ.*, vol. 238, 2020, Art. no. 111116.
- [15] R. Shang et al., "Near-real-time monitoring of land disturbance with harmonized Landsats 7–8 and Sentinel-2 data," *Remote Sens. Environ.*, vol. 278, 2022, Art. no. 113073.
- [16] Y. Wang, C. M. Albrecht, N. A. A. Braham, L. Mou, and X. X. Zhu, "Self-supervised learning in remote sensing: A review," *IEEE Geosci. Remote Sens. Mag.*, vol. 10, no. 4, pp. 213–247, Dec. 2022.
- [17] P. Jin, L. Mou, G.-S. Xia, and X. X. Zhu, "Anomaly detection in aerial videos with transformers," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5628213.
- [18] T. M. Tran, D. C. Bui, T. V. Nguyen, and K. Nguyen, "Transformer-based spatio-temporal unsupervised traffic anomaly detection in aerial videos," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 9, pp. 8292–8309, Sep. 2024.
- [19] D. Lu, P. Mausel, E. Brondizio, and E. Moran, "Change detection techniques," *Int. J. Remote Sens.*, vol. 25, no. 12, pp. 2365–2401, 2004.
- [20] W. Shi, M. Zhang, R. Zhang, S. Chen, and Z. Zhan, "Change detection based on artificial intelligence: State-of-the-art and challenges," *Remote Sens.*, vol. 12, no. 10, 2020, Art. no. 1688.
- [21] A. Shafique, G. Cao, Z. Khan, M. Asad, and M. Aslam, "Deep learning-based change detection in remote sensing images: A review," *Remote Sens.*, vol. 14, no. 4, 2022, Art. no. 871.
- [22] T. Bai et al., "Deep learning for change detection in remote sensing: A review," *Geo-spatial Inf. Sci.*, vol. 26, no. 3, pp. 262–288, 2023.
- [23] D. Peng, Y. Zhang, and H. Guan, "End-to-end change detection for high resolution satellite images using improved UNet++," *Remote Sens.*, vol. 11, no. 11, 2019, Art. no. 1382.
- [24] Z. Zheng, Y. Wan, Y. Zhang, S. Xiang, D. Peng, and B. Zhang, "CLNet: Cross-layer convolutional neural network for change detection in optical remote sensing imagery," *ISPRS J. Photogrammetry Remote Sens.*, vol. 175, pp. 247–267, 2021.
- [25] X. Wang et al., "Double U-Net (W-Net): A change detection network with two heads for remote sensing imagery," *Int. J. Appl. Earth Observ. Geoinformation*, vol. 122, 2023, Art. no. 103456.
- [26] R. C. Daudt, B. L. Saux, and A. Boulch, "Fully convolutional siamese networks for change detection," in *Proc. 25th IEEE Int. Conf. Image Process.*, IEEE, 2018, pp. 4063–4067.
- [27] D. Chicco, "Siamese neural networks: An overview," in *Artificial Neural Networks*. New York, NY, USA: Humana, 2021, pp. 73–94.
- [28] W. G. C. Bandara and V. M. Patel, "A transformer-based Siamese network for change detection," in *Proc. IGARSS IEEE Int. Geosci. Remote Sens. Symp.*, IEEE, 2022, pp. 207–210.
- [29] J. Chen et al., "DASNet: Dual attentive fully convolutional siamese networks for change detection in high-resolution satellite images," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 1194–1206, 2021.
- [30] H. Chen, Z. Qi, and Z. Shi, "Remote sensing image change detection with transformers," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2021, Art. no. 5607514.
- [31] W. Tang, K. Wu, Y. Zhang, and Y. Zhan, "A Siamese network based on multiple attention and multilayer transformers for change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5219015.
- [32] H. Yin et al., "Attention-guided Siamese networks for change detection in high resolution remote sensing images," *Int. J. Appl. Earth Observ. Geoinformation*, vol. 117, 2023, Art. no. 103206.
- [33] B. Zhao, H. Sui, and J. Liu, "Siam-DWNet: Flood inundation detection for SAR imagery using a cross-task transfer siamese network," *Int. J. Appl. Earth Observ. Geoinformation*, vol. 116, 2023, Art. no. 103132.
- [34] L. Yan and J. Jiang, "A hybrid siamese network with spatiotemporal enhancement and two-level feature fusion for remote sensing image change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 4403217.
- [35] Z.-H. You et al., "Crossed siamese vision graph neural network for remote sensing image change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 4409916.
- [36] Y. Zhou, J. Wang, J. Ding, B. Liu, N. Weng, and H. Xiao, "SIGNet: A siamese graph convolutional network for multi-class urban change detection," *Remote Sens.*, vol. 15, no. 9, 2023, Art. no. 2464.
- [37] M. Li, X. Liu, X. Wang, and P. Xiao, "Detecting building changes using multi-modal Siamese multi-task networks from very high resolution satellite images," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5404322.
- [38] Y. Yang, T. Chen, T. Lei, B. Du, A. K. Nandi, and A. Plaza, "SEDANet: A new Siamese ensemble difference attention network for building change detection in remotely sensed images," *IEEE Trans. Geosci. Remote Sens.*, vol. 62, 2024, Art. no. 4405316.

- [39] X. Niu, M. Gong, T. Zhan, and Y. Yang, "A conditional adversarial network for change detection in heterogeneous images," *IEEE Geosci. Remote Sens. Lett.*, vol. 16, no. 1, pp. 45–49, Jan. 2019.
- [40] Z. Du, X. Li, J. Miao, Y. Huang, H. Shen, and L. Zhang, "Concatenated deep learning framework for multi-task change detection of optical and SAR images," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 17, pp. 719–731, 2024.
- [41] C. Wu, B. Du, and L. Zhang, "Fully convolutional change detection framework with generative adversarial network for unsupervised, weakly supervised and regional supervised change detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 8, pp. 9774–9788, Aug. 2023.
- [42] H. Lyu, H. Lu, and L. Mou, "Learning a transferable change rule from a recurrent neural network for land cover change detection," *Remote Sens.*, vol. 8, no. 6, 2016, Art. no. 506.
- [43] L. Mou, L. Bruzzone, and X. X. Zhu, "Learning spectral-spatial-temporal features via a recurrent convolutional neural network for change detection in multispectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 2, pp. 924–935, Feb. 2019.
- [44] H. Chen, C. Wu, B. Du, L. Zhang, and L. Wang, "Change detection in multisource VHR images via deep Siamese convolutional multiple-layers recurrent neural network," *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 4, pp. 2848–2864, Apr. 2020.
- [45] U. Khushi, H. I. Dewangkoro, and A. M. Arymurthy, "Urban area change detection with combining CNN and RNN from Sentinel-2 multispectral remote sensing data," in *Proc. 3rd Int. Conf. Comput. Informat. Eng.*, IEEE, 2020, pp. 171–175.
- [46] M. Papadomanolaki, M. Vakalopoulou, and K. Karantzalos, "A deep multitask learning framework coupling semantic segmentation and fully convolutional LSTM networks for urban change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 9, pp. 7651–7668, Sep. 2021.
- [47] C. Shi, Z. Zhang, W. Zhang, C. Zhang, and Q. Xu, "Learning multiscale temporal-spatial-spectral features via a multipath convolutional LSTM neural network for change detection with hyperspectral images," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5529816.
- [48] T. Lei et al., "Ultralightweight spatial-spectral feature cooperation network for change detection in remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 4402114.
- [49] A. Codegoni, G. Lombardi, and A. Ferrari, "TINYCD: A (not so) deep learning model for change detection," *Neural Comput. Appl.*, vol. 35, no. 11, pp. 8471–8486, 2023.
- [50] L. Ding, K. Zhu, D. Peng, H. Tang, K. Yang, and L. Bruzzone, "Adapting segment anything model for change detection in VHR remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 62, 2024, Art. no. 5611711.
- [51] R. Touati, M. Mignotte, and M. Dahmane, "Anomaly feature learning for unsupervised change detection in heterogeneous images: A deep sparse residual model," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 13, pp. 588–600, 2020.
- [52] S. Saha, F. Bovolo, and L. Bruzzone, "Change detection in image time-series using unsupervised LSTM," *IEEE Geosci. Remote Sens. Lett.*, vol. 19, 2022, Art. no. 8005205.
- [53] B. Yang, L. Qin, J. Liu, and X. Liu, "UTRNet: An unsupervised time-distance-guided convolutional recurrent network for change detection in irregularly collected images," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 4410516.
- [54] Q. Xu, Y. Shi, J. Guo, C. Ouyang, and X. X. Zhu, "UCDFormer: Unsupervised change detection using a transformer-driven image translation," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5619917.
- [55] V. Ružička et al., "RaVEn: Unsupervised change detection of extreme events using ml on-board satellites," *Sci. Rep.*, vol. 12, no. 1, 2022, Art. no. 16939.
- [56] R. Yadav, A. Nascetti, H. Azizpour, and Y. Ban, "Unsupervised flood detection on SAR time series using variational autoencoder," *Int. J. Appl. Earth Observ. Geoinformation*, vol. 126, 2024, Art. no. 103635.
- [57] P. S. White, "Natural disturbance and patch dynamics: An introduction," in *The Ecology Natural Disturbance Patch Dynamics*. New York, NY, USA: Academic Press, Inc. 1985, pp. 3–13.
- [58] M. G. Turner, "Disturbance and landscape dynamics in a changing world," *Ecology*, vol. 91, no. 10, pp. 2833–2849, 2010.
- [59] J. Verbesselt, R. Hyndman, G. Newnham, and D. Culvenor, "Detecting trend and seasonal changes in satellite image time series," *Remote Sens. Environ.*, vol. 114, no. 1, pp. 106–115, 2010.
- [60] 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 Sens. Environ.*, vol. 114, no. 12, pp. 2897–2910, 2010.
- [61] J. Verbesselt, A. Zeileis, and M. Herold, "Near real-time disturbance detection using satellite image time series," *Remote Sens. Environ.*, vol. 123, pp. 98–108, 2012.
- [62] Z. Zhu and C. E. Woodcock, "Continuous change detection and classification of land cover using all available landsat data," *Remote Sens. Environ.*, vol. 144, pp. 152–171, 2014.
- [63] 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 Sens. Environ.*, vol. 252, 2021, Art. no. 112167.
- [64] E. Hamunyela, J. Verbesselt, and M. Herold, "Using spatial context to improve early detection of deforestation from Landsat time series," *Remote Sens. Environ.*, vol. 172, pp. 126–138, 2016.
- [65] J. Li, Z.-L. Li, H. Wu, and N. You, "Trend, seasonality, and abrupt change detection method for land surface temperature time-series analysis: Evaluation and improvement," *Remote Sens. Environ.*, vol. 280, 2022, Art. no. 113222.
- [66] M. Schwieder et al., "Mapping grassland mowing events across Germany based on combined Sentinel-2 and Landsat 8 time series," *Remote Sens. Environ.*, vol. 269, 2022, Art. no. 112795.
- [67] X. Yang, Z. Zhu, S. Qiu, K. D. Kroeger, Z. Zhu, and S. Covington, "Detection and characterization of coastal tidal wetland change in the northeastern US using landsat time series," *Remote Sens. Environ.*, vol. 276, 2022, Art. no. 113047.
- [68] V. Fotakidis, I. Chrysafis, G. Mallinis, and N. Koutsias, "Continuous burned area monitoring using bi-temporal spectral index time series analysis," *Int. J. Appl. Earth Observ. Geoinformation*, vol. 125, 2023, Art. no. 103547.
- [69] S. Ye, Z. Zhu, and G. Cao, "Object-based continuous monitoring of land disturbances from dense landsat time series," *Remote Sens. Environ.*, vol. 287, 2023, Art. no. 113462.
- [70] M. Hussain, D. Chen, A. Cheng, H. Wei, and D. Stanley, "Change detection from remotely sensed images: From pixel-based to object-based approaches," *ISPRS J. Photogrammetry Remote Sens.*, vol. 80, pp. 91–106, 2013.
- [71] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 9729–9738.
- [72] I. Misra and L. v. d. Maaten, "Self-supervised learning of pretext-invariant representations," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 6707–6717.
- [73] D. Hong et al., "SpectralGPT: Spectral remote sensing foundation model," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 8, pp. 5227–5244, Aug. 2024.
- [74] F. Liu et al., "RemoteCLIP: A vision language foundation model for remote sensing," *IEEE Trans. Geosci. Remote Sens.*, vol. 62, 2024, Art. no. 5622216.
- [75] X. X. Zhu et al., "On the foundations of earth and climate foundation models," 2024, *arXiv:2405.04285*.
- [76] W. Liu, W. Luo, D. Lian, and S. Gao, "Future frame prediction for anomaly detection—A new baseline," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 6536–6545.
- [77] Y. Chang, Z. Tu, W. Xie, and J. Yuan, "Clustering driven deep autoencoder for video anomaly detection," in *Proc. 16th Eur. Conf., Comput. Vis.*, Glasgow, U.K., Springer, Aug. 2020, pp. 329–345.
- [78] H. Park, J. Noh, and B. Ham, "Learning memory-guided normality for anomaly detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 14372–14381.
- [79] Y. Zhang, X. Nie, R. He, M. Chen, and Y. Yin, "Normality learning in multispace for video anomaly detection," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 31, no. 9, pp. 3694–3706, Sep. 2021.
- [80] C. Huang, Z. Wu, J. Wen, Y. Xu, Q. Jiang, and Y. Wang, "Abnormal event detection using deep contrastive learning for intelligent video surveillance system," *IEEE Trans. Ind. Informat.*, vol. 18, no. 8, pp. 5171–5179, Aug. 2022.
- [81] Z. Liu, Y. Nie, C. Long, Q. Zhang, and G. Li, "A hybrid video anomaly detection framework via memory-augmented flow reconstruction and flow-guided frame prediction," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 13588–13597.
- [82] G. Yu, S. Wang, Z. Cai, X. Liu, C. Xu, and C. Wu, "Deep anomaly discovery from unlabeled videos via normality advantage and self-paced refinement," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 13987–13998.
- [83] D. Gong et al., "Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 1705–1714.

- [84] S. Akcay, A. Atapour-Abarghouei, and T. P. Breckon, "GANomaly: Semi-supervised anomaly detection via adversarial training," in *Proc. 14th Asian Conf. Comput. Vis.*, Perth, Australia, Springer, Dec. 2018, pp. 622–637.
- [85] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [86] G. Sharir, A. Noy, and L. Zelnik-Manor, "An image is worth 16x16 words, what is a video worth?," 2021, *arXiv:2103.13915*.
- [87] Y. Yuan and L. Lin, "Self-supervised pretraining of transformers for satellite image time series classification," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 474–487, 2021.
- [88] Y. Yuan, L. Lin, Q. Liu, R. Hang, and Z.-G. Zhou, "SITS-Former: A pre-trained spatio-spectral-temporal representation model for Sentinel-2 time series classification," *Int. J. Appl. Earth Observ. Geoinformation*, vol. 106, 2022, Art. no. 102651.
- [89] I. Dumeur, S. Valero, and J. Inglada, "Self-supervised spatio-temporal representation learning of satellite image time series," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 17, pp. 4350–4367, 2024.
- [90] A. Vaswani, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 6000–6010.
- [91] G. Tseng, R. Cartuyvels, I. Zvonkov, M. Purohit, D. Rolnick, and H. Kerner, "Lightweight, pre-trained transformers for remote sensing timeseries," 2023, *arXiv:2304.14065*.
- [92] EO Research, "How to normalize satellite images for deep learning," *Medium*, 2022. [Online]. Available: <https://medium.com/sentinel-hub/how-to-normalize-satellite-images-for-deep-learning-d5b668c885af>
- [93] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. 18th Int. Conf., Med. Image Comput., Comput.-Assist. Interv.*, Munich, Germany, Springer, Oct. 2015, pp. 234–241.
- [94] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [95] National Water and Climate Center, "Air and water database public reports—leavitt meadows," *Natural Resour. Conservation Serv.*, 2024. Accessed: Aug. 27, 2024. [Online]. Available: <https://wcc.sc.gov.usda.gov/nwcc/site?sitenum=575>
- [96] N. I. Gasparri and H. R. Grau, "Deforestation and fragmentation of Chaco dry forest in NW Argentina (1972–2007)," *Forest Ecol. Manage.*, vol. 258, no. 6, pp. 913–921, 2009.
- [97] V. Fehlenberg, M. Baumann, N. I. Gasparri, M. Piquer-Rodriguez, G. Gavier-Pizarro, and T. Kuemmerle, "The role of soybean production as an underlying driver of deforestation in the South American Chaco," *Glob. Environ. Change*, vol. 45, pp. 24–34, 2017.
- [98] WaterNSW, "Copeton dam," (n.d.). Accessed: Dec. 09, 2024. [Online]. Available: <https://www.watersnsw.com.au/nsw-dams/regional-nsw-dams/copeton-dam>
- [99] Department of Planning and Environment, "Gwydir valley annual surface water quality report 2021–2022," NSW Government, Tech Rep. DOC22/291568, 2022.
- [100] S. Mishra and D. R. Mishra, "Normalized difference chlorophyll index: A novel model for remote estimation of chlorophyll-a concentration in turbid productive waters," *Remote Sens. Environ.*, vol. 117, pp. 394–406, 2012.
- [101] WaterNSW, "WaterInsights," (n.d.). Accessed: Dec. 09, 2024. [Online]. Available: <https://waterinsights.watersnsw.com.au/11985-gwydir-regulated-river/storage>
- [102] Lismore City Council, "History of lismore flood events 1870–2022," Lismore City Council, Lismore, NSW, Australia, Jun. 2022.
- [103] Lismore City Council, "Lismore flood response report," Lismore, NSW, Australia, Apr. 2022.
- [104] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2013, *arXiv:1312.6114*.
- [105] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, vol. 28, pp. 3483–3491.
- [106] A. Radoi, "Multimodal satellite image time series analysis using GAN-based domain translation and matrix profile," *Remote Sens.*, vol. 14, no. 15, 2022, Art. no. 3734.
- [107] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM Comput. Surv.*, vol. 53, no. 3, pp. 1–34, 2020.
- [108] X. Yang, Z. Song, I. King, and Z. Xu, "A survey on deep semi-supervised learning," *IEEE Trans. Knowl. Data Eng.*, vol. 35, no. 9, pp. 8934–8954, Sep. 2023.
- [109] F. Pourpanah et al., "A review of generalized zero-shot learning methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 4, pp. 4051–4070, Apr. 2023.
- [110] S. Garske, B. Evans, C. Artlett, and K. Wong, "ERX: A fast real-time anomaly detection algorithm for hyperspectral line scanning," *IEEE Trans. Geosci. Remote Sens.*, vol. 63, 2025, Art. no. 5503617.



Samuel Garske (Graduate Student Member, IEEE) received the dual bachelor's degree (Hons.) in mechanical engineering, the bachelor of commerce degree in business analytics, and the Ph.D. degree in engineering from The University of Sydney, Sydney, NSW, Australia, in 2019 and 2025.

He was a Visiting Student Researcher with the National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory, Pasadena, CA, USA. He has also been a Visiting Research Fellow with the Future Lab of Artificial Intelligence for Earth

Observation in Munich, Germany; a joint collaboration between the Technical University of Munich (TUM) and the German Aerospace Center (DLR). He was also a student with the Australian Research Council Training Center for CubeSats, UAVs, and Their Applications (CUAVA). He is currently a Research Associate with the University of Sydney. His research interests include computer vision, anomaly detection, object detection, remote sensing, and applied machine learning for Earth observation applications.



Konrad Heidler received the bachelor's degree in mathematics, the master's degree in mathematics in data science, and the Doctorate in Engineering degree (Dr.-Ing.) from the Technical University of Munich (TUM), Munich, Germany, in 2017, 2020, and 2024, respectively.

He was a Postdoctoral Researcher with TUM, where he lead the working group for visual learning and reasoning at the Chair for Data Science in Earth Observation. He is now a Senior Data Scientist with Hula Earth, Munich, Germany. His research interests

focus on the application of deep learning for remote sensing in polar regions, solving reasoning tasks with deep learning, and applications of self- and semi-supervised learning in Earth observation.



Bradley Evans received the dual bachelor of science degree (Hons.) in environmental science, the bachelor of science degree in energy studies, and the Ph.D. degree in environmental science from Murdoch University, Murdoch, WA, Australia, in 2009 and 2013.

He was a Director of Australia's Terrestrial Ecosystem Research Network, Ecosystem Modeling and Scaling Infrastructure, at Macquarie University and the University of Sydney. He was with the Australian Government Department of Defence, where he served

on the AquaWatch and other committees related to Earth Observation. He was a Director of Sydney Informatics Hub during his time at the University of Sydney. Since 2023, he is currently an Associate Professor with the School of Environmental and Rural Science, University of New England, Armidale, NSW, Australia. During his tenure he has established the Earth Observation Laboratory with a special focus on water and wildlife habitat (koalas), and riverine water quality. He has worked on projects associated with NASA's OCO2 mission, working to find new ways of modeling plant growth using plant fluorescence. As a Chief Investigator of the Australian Research Council Training Centre on CubeSats, UAVs and Their Applications (CUAVA), Bradley has led the development of an OpenSource Hyperspectral Imaging Spectrometer, OpenHSI in collaboration with the Rochester Institute of Technology. He is also a community contributor to NASA JPL's Surface Biology and Geology Study. His research interests include environmental science, biodiversity, and applying hyperspectral imaging spectroscopy to solve real world problems.



Kee Choon Wong received the bachelor's degree in aeronautical engineering and the Ph.D. degree in aerospace engineering from the University of Sydney, Sydney, NSW, Australia, in 1985 and 1994.

Since 1992, he has been an active full-time academic in aerospace engineering with the School of Aerospace Mechanical and Mechatronic Engineering, the University of Sydney. He is currently a Professor of aerospace engineering, Director for Aeronautical Engineering, and the Deputy Head of School. Since 1988, he is a pioneering uncrewed aircraft systems (UAS) Researcher, having worked on multidisciplinary and morphing airframe design, instrumentation, control, system integration, applications, flight testing, and project management. He is very active in facilitating the next generation of UAS and aerospace engineers through his introduction of unique experiential learning and global engineering design opportunities. He leads a small UAS research team and has international R&D collaborations. UAS designed and developed by his group have been used in several collaborative research projects with industry. His research interests include CubeSats, UAVs, and their applications (CUAVA).

Prof. Wong was the founding President of the industry-focused Australian Association for Uncrewed Systems (AAUS) and served in that role for seven years until 2015. He is an Associate Fellow of the American Institute of Aeronautics and Astronautics (AIAA), and a Fellow of the Royal Aeronautical Society (RaeS). He is a Chief Investigator with the Australian Research Council Training Centre.



Xiao Xiang Zhu (Fellow, IEEE) received the master's (M.Sc.) degree, the doctor of engineering (Dr.-Ing.) degree, and the "Habilitation" degree in the field of signal processing from the Technical University of Munich (TUM), Munich, Germany, in 2008, 2011, and 2013, respectively.

She is the Chair Professor in data science in Earth observation with Technical University of Munich (TUM), Munich, Germany, and was the founding Head of the Department "EO Data Science" with the Remote Sensing Technology Institute, German Aerospace Center (DLR). Since May 2020, she has been the PI and Director of the international future AI lab "AI4EO—Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond," Munich, Germany. Since October 2020, she has also been a Director of the Munich Data Science Institute (MDSI), TUM. From 2019 to 2022, Zhu has been a coCoordinator of the Munich Data Science Research School (www.mu-ds.de) and the Head of the Helmholtz Artificial Intelligence – Research Field "Aeronautics, Space and Transport." She was a Guest Scientist or Visiting Professor with the Italian National Research Council (CNR-IREA), Naples, Italy, Fudan University, Shanghai, China, the University of Tokyo, Tokyo, Japan, and the University of California, Los Angeles, CA, USA, in 2009, 2014, 2015, and 2016, respectively. She is currently a Visiting AI Professor with ESA's Phi-lab, Frascati, Italy. Her main research interests include remote sensing and Earth observation, signal processing, machine learning and data science, with their applications in tackling societal grand challenges, e.g., global urbanization, UN's SDGs, and climate change.

Dr. Zhu has been a Member of young academy (Junge Akademie/Junges Kolleg), the Berlin-Brandenburg Academy of Sciences and Humanities and the German National Academy of Sciences Leopoldina and the Bavarian Academy of Sciences and Humanities. She is a Fellow of the Academia Europaea (the Academy of Europe). She is in the scientific advisory board in several research organizations, among others the German Research Center for Geosciences (GFZ, 2020–2023) and Potsdam Institute for Climate Impact Research (PIK). She is an Associate Editor of IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING, *Pattern Recognition*, and served as the area Editor responsible for special issues of IEEE SIGNAL PROCESSING MAGAZINE (2021–2023). She is a Fellow of AIAA and ELLIS.