
Unsupervised Change Detection of Extreme Events Using ML On-Board

Vít Ružička^{1*}, Anna Vaughan², Daniele De Martini¹, James Fulton³,
Valentina Salvatelli⁴, Chris Bridges⁵, Gonzalo Mateo-Garcia⁶, Valentina Zantedeschi⁷
¹University of Oxford, ²University of Cambridge, ³University of Edinburgh, ⁴Microsoft,
⁵University of Surrey, ⁶University of Valencia, ⁷INRIA, University College London

Abstract

In this paper, we introduce **RaVÆn**, a lightweight, unsupervised approach for change detection in satellite data based on Variational Auto-Encoders (VAEs) with the specific purpose of on-board deployment. Applications such as disaster management enormously benefit from the rapid availability of satellite observations. Traditionally, data analysis is performed on the ground after all data is transferred – downlinked – to a ground station. Constraint on the downlink capabilities therefore affects any downstream application. In contrast, **RaVÆn** pre-processes the sampled data directly on the satellite and flags changed areas to prioritise for downlink, shortening the response time. We verified the efficacy of our system on a dataset composed of time series of catastrophic events – which we plan to release alongside this publication – demonstrating that **RaVÆn** outperforms pixel-wise baselines. Finally we tested our approach on resource-limited hardware for assessing computational and memory limitations.

1 Introduction

Satellite observations of the Earth’s surface provide a vital data source for diverse environmental applications, including disaster management, landcover change detection, and ecological monitoring. Currently, sensors on satellites collect and downlink all data points for further processing on the ground; yet, limitations in downlink capacity and speed result in delayed data availability and inefficient use of ground stations. These limitations affect time-sensitive applications such as disaster management where data are required with as low latency as possible to inform decision making in real time. This problem is set to worsen as the sensing resolution and number of satellites in orbit increase together with further restrictions on radio-frequency spectrum sharing licensing.

A solution is to apply processing on-board to identify the most useful data for a particular scenario, and prioritise this for rapid downlink. There has been considerable recent interest in using machine learning on-board for this processing. Existing work has focused on the deployment of supervised classifiers for applications such as identifying clouds [1] or floods [2]. Supervised learning has a significant drawback: only events of a particular type determined at training time will be flagged and prioritised for downlink, with no generalisation to new event types, imager specifications, and lighting or local features.

In this work we propose a new fully unsupervised novelty-detection model suitable for deployment on remote sensing platforms. We use a *Variational AutoEncoder* (VAE) [3] to generate a latent representation of incoming sensor data over a particular region. A novelty score is assigned to this data using the distance in the latent space between representations from consecutive passes. This offers a substantial advantage over existing supervised methods as any change between passes can be detected on-board at inference time, regardless of the availability of training data.

*Correspondence to: vit.ruzicka@cs.ox.ac.uk

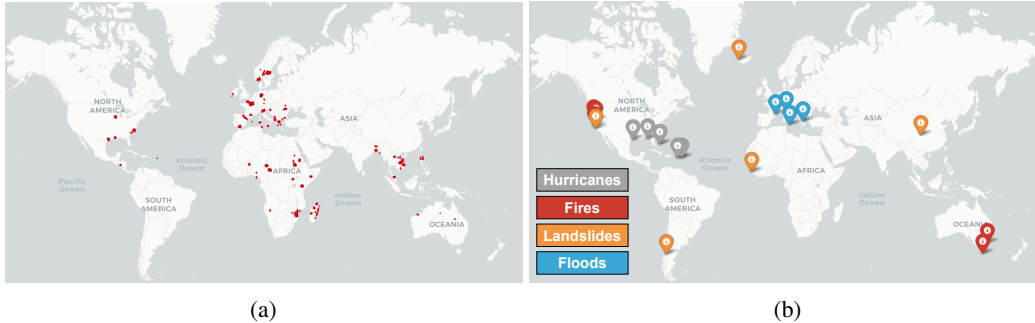


Figure 1: Locations used for training (a) and validation (b) images.

We evaluate the performance of this model detecting changes in land-surface observations from the Sentinel-2 Multispectral Instrument on time series of natural disaster events where obtaining data with low-latency is of importance, including floods, landslides, wildfires and hurricanes. The unsupervised novelty detection model is demonstrated to assign higher novelty scores to regions of known change and outperforms computer-vision image differencing baselines. We further demonstrate via experiments on constrained Xilinx Pynq hardware that this model is suitable for deployment on a remote sensing platform.

1.1 Application Context

Our proposed method **RaVÆn** facilitates intelligent decision making on board of a satellite for rapid disaster response. With this application in mind it has been tested on constrained hardware to similar to that available on remote sensing platforms. Using this system we can detect extreme events and rapidly prioritize data over the affected regions for downlink. Vital data will therefore be available with lower latency to inform decision making during rapidly evolving events. In the rest of this section we frame our work in the context of machine learning literature.

Anomaly detection The use of VAEs has been explored for unsupervised anomaly detection in [4], where the model reconstruction error has been used as anomaly score. Our approach differs in that, instead of basing our predictions on the reconstruction error of a single input, which has been shown in [5] to be an unreliable indicator in the unsupervised context, we consider a sequence of input images from the same location. Our problem is better framed as change detection.

Change detection The need for annotations of supervised change detection techniques, such as siamese networks in [6], can be reduced using active learning approaches as demonstrated in [7], but then it still lacks in terms of generality. The main challenge of unsupervised change detection is being able to distinguish changes of interest from spurious change due to noise. Many existing approaches [8, 9, 10] achieve this by combining dimensionality reduction techniques, such as Principal Component Analysis, and clustering, such as k -means, to detect only relevant change between images of consecutive passes. Approaches based on neural networks (see [11] for a review) rely instead on supervised auxiliary tasks, such as semantic segmentation, to extract informative features that are then used to detect change in a time series. Our method leverages neural networks without requiring supervision at any stage.

2 Data

As part of this study, we compile a new dataset to evaluate the proposed unsupervised change detection models. Images are taken from the Sentinel-2 multi-spectral imager (MSI) instrument² and include the ten highest resolution channels with all channels interpolated to the highest resolution of 10m. Training data are taken from the *WorldFloods* dataset of [2] locations (Figure 1a), with a total of 233 scenes and a time series of five images per scene.

²<https://sentinels.copernicus.eu/web/sentinel/user-guides/sentinel-2-msi/overview>




Figure 2: Example of validation sample – in this case, a hurricane event – and its corresponding ground-truth mask (which contains labels of change and clouds).

The validation set, consists of Sentinel-2 time series for four classes of disasters: hurricanes, fire burn scars, landslides and floods (Figure 1b). We identified events in each of these classes through an extensive search of Sentinel-2 records aided by the Copernicus EMS system³. Each event in the validation set consists of a time series of five images where the first four images are taken before the disaster while the fifth image is taken afterwards. To mitigate the effects of cloud cover, we discarded validation images with greater than 20% cloud cover. Events are only included where all images are within 180 days before and 90 days after the event. For each event a change mask was hand annotated to mark corresponding to the difference between the final two images in the time series, for example Figure 2. We emphasize that these labels are used for evaluation only. The distribution of data in the validation dataset is further detailed in Appendix A.3.

3 Methodology

Preprocessing Tiles of 32x32 pixels are extracted from the Sentinel-2 scenes described above and used as inputs to the considered models. These are further normalized by applying a log transform and scaling to constrain them to the $[0, 1]$ interval.

Model We use a VAE (illustrated in Figure 6 of the appendix) to learn an meaningful embedding space for change detection. The VAE is chosen as it guarantees a meaningful distance metric on the embedding space. The model is trained to reconstruct the original input tile from its compressed representation distribution in the embedding space. Once trained we use this model as a feature extractor to encode individual tiles, e.g. $x_t^{a,b}$ and $x_{t-1}^{a,b}$, and compare their latent representations to measure the change between time-steps t and $t - 1$. Parameters a, b correspond to the indexed location in the original image. Comparing compressed representations of images improves robustness to noise and reduces the computational and memory requirements of storing images from previous passes, which is a critical in a constrained environment. We denote by $\mathcal{N}_t^{a,b} = \mathcal{N}(\mu_t^{a,b}, \sigma_t^{a,b})$ the encoded Gaussian distribution of tile $x_t^{a,b}$ with $\mu_t^{a,b}, \sigma_t^{a,b} \in \mathbb{R}^n$. In the following, we fix the latent size to $n = 128$ after the initial experiments indicated that larger latent sizes did not yield improved results.

Change detection novelty score We compute a change score for tiles $x_{t-1}^{a,b}, x_t^{a,b}$ using either the KL divergence between encoded distributions $\mathcal{N}_{t-1}^{a,b}$ and $\mathcal{N}_t^{a,b}$, or the Euclidean or cosine distance between mean encoded vectors $\mu_{t-1}^{a,b}$ and $\mu_t^{a,b}$, that we generally note as $D_{\text{VAE}}(x_{t-1}^{a,b}, x_t^{a,b})$. We extend the methodology to use a longer time series of k past images by retaining the minimal distance as:

$$S_{\text{VAE}}(x_t^{a,b}) = \min_{i=1..k} D_{\text{VAE}}(x_{t-i}^{a,b}, x_t^{a,b}) \quad (1)$$

To compare the performance of this approach to simpler on-board processing methods that does not make use of machine learning, we compare our method to a baseline which compares tiles

³<https://emergency.copernicus.eu/mapping/list-of-activations-rapid>




Figure 3: Example of **RaVAEn**, the suggested decision making process on board of the satellite (for a time window of $k = 1$).

directly in the input space using the Euclidean or the cosine distance and after applying the same data pre-processing as for the VAE.

The proposed workflow outlined in this section is summarised in Figure 3.

4 Experimental Setting and Reproducibility

We use different environments for training the VAE and for inference. For development (training and validation of the model), we use a `n1-standard-16` instance on Google Cloud Platform with two NVIDIA Tesla V100 GPUs. In addition we measure the performance of the models on the Xilinx Pynq FPGA board with limited compute power, 650MHz ARM Cortex-A9 CPU and 512MB RAM which emulates the resources available on a typical small satellite. We report the metrics for the methodology deployed on real hardware in Appendix A.2.

Model architecture design choices To minimize the size of the model and maximize efficiency on constrained devices we conducted a parametric search over both the number of layers and number of units per layer in both the encoder network E and the decoder network D . More precisely we tested three different model architecture configurations (*small*, *medium* and *large*). Details and results of this experiment are included in Appendix A.2, where we show their maintained performance while reaching fast processing speeds of only ~ 2 seconds on the Xilinx Pynq FPGA board. In the following, we report results for the *large* VAE configuration.

5 Results

Figure 4 shows a qualitative comparison between the VAE model developed in this study and the image differencing baseline. The *before* image shows a river that floods and therefore changes colour in the *after* image. The labels and the change scores from our *large* VAE and baseline methods are shown in the bottom row of the figure. In this case, the scores were calculated using a history of $k = 3$ frames, although only the most recent *before* frame is shown for brevity. In this example, our method – the cosine embedding – produces a change map that is crisper than the cosine baseline; for instance, the small flooded canal can be seen in the cosine embedding image but not in the baseline. In a similar fashion, Figure 5 shows a qualitative comparison in the case of a burnt-area detection.

The change-score maps, like those in Figures 4 and 5, were produced for every image in the evaluation set. We use these maps and our labels to calculate the area under the precision-recall curve. We produce the curve tile-wise, so that each individual tile across each image is treated as a positive or negative example of change, rather than treating the full image as one example. This means our quality metric is sensitive to the fact that our evaluation images are not equal; they have different number of tiles and different ratios of positive pixels (as reported in Table 5). We also ignore tiles that have clouds in the *after* image or in the most recent image before the event. We produce a precision-recall curve for each of the four different event types in our evaluation set.

Table 1 reports the results of our change detection experiments for all disaster types. We found that cosine distance, applied on the input space or on the embeddings, generally provides the best scores. Surprisingly KL-divergence is the lowest-performing metric, and is beaten by both cosine and Euclidean embedding scores in all events, even though these methods do not use the variance values calculated by the VAE. Metrics based on the VAE embedding outperforms the baseline on floods, hurricanes and fires, and reaches similar performance on landslides.

Table 2 shows the effects of including a longer frame history. When three previous images are provided instead of just one, both the embedding and baseline perform better except for the case of




Figure 4: Comparison of the change detected using the baseline and the *large* VAE method on an example of a flooding river. Two images immediately before and immediately after a change are shown, along with the human labels of change and the calculated change scores. Both methods used a history of $k = 3$ frames.



Figure 5: Additional comparison of the change detected using the baseline and the *large* VAE method on an example of a fire disaster. Both methods used a history of $k = 3$ frames. The cosine baseline prediction seems to more closely copy the details present in the image, making it susceptible to small, noisy variations between the two images.

landslide dataset where the cosine baseline with memory 1 beats memory 3 with a small margin. The table also shows that our method of detecting significant change based on the embedding space outperforms the baseline in every dataset when $k = 3$.

Table 1: Area under the precision-recall curve for baseline and VAE methods with time window $k = 1$ (averaged over 5 runs).

Detection method	Dataset			
	Landslides	Floods	Hurricanes	Fires
Cosine baseline	0.629	0.378	0.513	0.818
Euclidean baseline	0.267	0.326	0.351	0.770
Cosine embedding	0.599 ± 0.012	0.448 ± 0.011	0.676 ± 0.014	0.833 ± 0.008
Euclidean embedding	0.266 ± 0.004	0.450 ± 0.007	0.478 ± 0.019	0.800 ± 0.011
KL-Divergence	0.258 ± 0.022	0.247 ± 0.018	0.301 ± 0.035	0.731 ± 0.016

Table 2: Area under the precision-recall curve for the best performing metrics from table 1 with and without an extended history k (averaged over 5 runs).

Detection method	k	Dataset			
		Landslides	Floods	Hurricanes	Fires
Cosine baseline	1	0.629	0.378	0.513	0.818
	3	0.622	0.378	0.570	0.865
Cosine embedding	1	0.599 ± 0.012	0.448 ± 0.011	0.676 ± 0.014	0.833 ± 0.008
	3	0.759 ± 0.024	0.443 ± 0.009	0.726 ± 0.011	0.913 ± 0.008

6 Conclusion

In conclusion, we introduce a new method **RaVÆn** for unsupervised change detection in remote sensing data using a VAE. Our method is evaluated on a new dataset of remote sensing images of disasters which we aim to release with our work. The proposed model outperforms a classical computer vision baseline in all of the tested disaster classes. This demonstrates that **RaVÆn** is a robust change detection method and suitable for application in improving data acquisition for disaster response.

Acknowledgments and Disclosure of Funding

This work has been enabled by Frontier Development Lab (FDL) Europe, a public partnership between the European Space Agency (ESA) Phi-Lab (ESRIN) and ESA Mission Operations (ESOC), Trillium Technologies and the University of Oxford; the project has been also supported by Google Cloud, D-Orbit and Planet. The authors would like to thank all FDL faculty members, Atılım Güneş Baydin and Yarin Gal (University of Oxford), Chedy Raissi (INRIA), Brad Neuberger (Planet) and Nicolas Longép  (ESA ESRIN) for discussions and comments throughout the development of this work.

References

- [1] Gianluca Giuffrida, Lorenzo Diana, Francesco de Gioia, Gionata Benelli, Gabriele Meoni, Massimiliano Donati, and Luca Fanucci. CloudScout: A Deep Neural Network for On-Board Cloud Detection on Hyperspectral Images. *Remote Sensing*, 12(14):2205, January 2020.
- [2] Gonzalo Mateo-Garcia, Joshua Veitch-Michaelis, Lewis Smith, Silviu Vlad Oprea, Guy Schumann, Yarin Gal, Atılım Güneş Baydin, and Dietmar Backes. Towards global flood mapping onboard low cost satellites with machine learning. *Scientific Reports*, 11(1):7249, March 2021. ISSN 2045-2322.
- [3] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2013.
- [4] Daniel Angerhausen, Valentin Tertius Bickel, and Lesnikowski Adam. Unsupervised distribution learning for lunar surface technosignature detection. *Earth and Space Science Open Archive*, page 1, 2020.
- [5] Nicholas Merrill and Azim Eskandarian. Modified autoencoder training and scoring for robust unsupervised anomaly detection in deep learning. *IEEE Access*, 2020.
- [6] Rodrigo Caye Daudt, Bertr Le Saux, and Alexandre Boulch. Fully convolutional siamese networks for change detection. *2018 25th IEEE International Conference on Image Processing (ICIP)*, Oct 2018.
- [7] Vít Ruřička, Stefano D’Aronco, Jan D Wegner, and Konrad Schindler. Deep active learning in remote sensing for data efficient change detection. In *Proceedings of MACLEAN: MACHine Learning for EArth ObservatiON Workshop (ECML/PKDD 2020)*, volume 2766, 2020.
- [8] Turgay Celik. Unsupervised change detection in satellite images using principal component analysis and k -means clustering. *IEEE Geoscience and Remote Sensing Letters*, 2009.
- [9] Turgay Çelik and Chandra V. Curtis. Resolution selective change detection in satellite images. In *ICASSP*, 2010.
- [10] Yongqiang Cheng, Heng-Chao Li, Turgay Çelik, and Fan Zhang. Frft-based improved algorithm of unsupervised change detection in SAR images via PCA and k -means clustering. In *IGARSS*, 2013.
- [11] Kevin Louis de Jong and Anna Sergeevna Bosman. Unsupervised change detection in satellite images using convolutional neural networks. In *IJCNN*. IEEE, 2019.
- [12] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 1(10):e3, 2016. URL <http://distill.pub/2016/deconv-checkerboard>. Accessed on 1.9.2021.
- [13] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8695–8704, 2020.

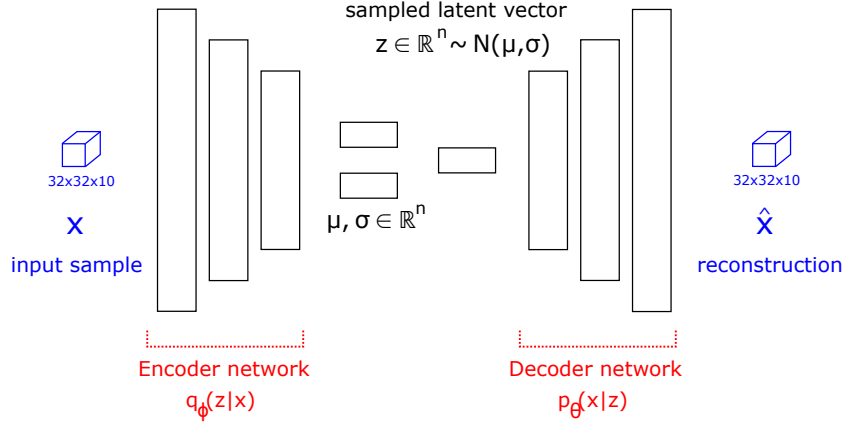


Figure 6: Illustration of the VAE model, n denotes the latent dimension. Number and type of the used layers can be adjusted given the limitations of the used hardware as is explored in more detail in A.2.

A Appendix

A.1 Model architecture

The encoder of our VAE was composed of a series of downsampling blocks. Each downsampling block first had a 2D convolutional layer with kernel size 3, stride size 2, and zero padding of 1, such that the dimensions are halved in the spatial domain. Following this layer, the block also had a sequence of extra 2D convolutional layers (the number *extra depth* referred to in table 4). Skip connections were used so that the *extra depth* convolutional layers formed a residual block. The network could then easily learn to skip these non-downsampling layers. In the residual block, the number of hidden channels and image size were conserved. Each convolution layer used leaky ReLU activations and batch normalisation. Following a given number of downsampling blocks, the result was flattened and further reduced in dimension using a fully connected layer which output the mean and log variance. The decoder was essentially the encoder in reverse. The upsampling method used was nearest neighbour upsampling followed by a single convolution. This method was preferred over transpose convolution to avoid checkerboard artefacts [12, 13].

The main model presented in this paper is denoted as *large* on Table 4, it used 3 downsampling blocks with 32, 64, 128 channels on each successively smaller scale. The network was further squeezed via the fully connected layer to a latent dimension of 128. We used extra depth value of 2, so after each downscale convolution there was a residual block of 2 additional convolutional layers.

A.2 Model timings

Table 3: Area under the precision recall curve and timings for different sizes of model (averaged over 5 runs). The AUPRC results are for the cosine similarity of the embedding with a history of 3 frames.

	Dataset			
	Landslides	Floods	Hurricanes	Fires
Small model	0.748 \pm 0.014	0.445 \pm 0.014	0.748 \pm 0.002	0.907 \pm 0.002
Medium model	0.758 \pm 0.007	0.428 \pm 0.004	0.738 \pm 0.018	0.912 \pm 0.003
Large model	0.759 \pm 0.024	0.443 \pm 0.009	0.726 \pm 0.011	0.913 \pm 0.001

It is strongly desired that this change detection method could be run onboard a satellite, so that images can be filtered or prioritised before they reach the ground. Therefore the time it takes to run change detection is of the utmost importance. We need to design our models so that the low power hardware, available onboard a mission, can keep up with the incoming stream of data. We therefore tested the accuracy of our model with different numbers of channels, and different numbers of layers.

Table 3 shows the accuracy of a few variations of model size and the time it took to process a 574×509 px image (approx. $5\text{km} \times 5\text{km}$ at Sentinel-2 10m resolution) whilst running on the CPU of a Xilinx PYNQ. We see the the results of all tested models are comparable and that it is reasonable to aim for the smallest model, which takes only 2.06 seconds to process the patch. Running onboard the PYNQ means that there is considerable potential to speed up this runtime by a large factor by deploying directly on the FPGA module rather than using the board’s CPU.

Table 4: Differences in the architecture for different models of table 3. Note that during inference, we only need the encoder network of the VAE model. We also only need to process the newly acquired image to obtain their latent representation, while the latent vectors of the previous image can be loaded.

	Total parameters (millions)	Encoder parameters (millions)	Runtime (seconds)	Extra depth	Hidden channels
Small model	0.443	0.285	2.06	0	16, 32, 64
Medium model	0.979	0.617	4.86	0	32, 64, 128
Large model	1.500	1.005	13.98	2	32, 64, 128

A.3 Dataset description

We describe the statistics of the manually annotated validation dataset in Table 5. While each type of the event is represented by similar amount of locations, the affected area varies dramatically with different disasters. Namely the area of burn scars in the *Fire* dataset has both the largest area of effect and also the largest proportion of changed pixels to all not cloudy pixels (reported as positive ratio).

Table 5: Validation dataset details. Each location is captured in 4 time-steps before the event and once after the event (it is however counted only once into the cumulative square km). Note that only the last pair of images is labelled with changed and not changed pixel map - and that the reported percentage only considers this map.

	Number of locations	Cumulative area (km ²)	Positive rate (%)
Landslides	5	108	10.48
Floods	4	1301	6.74
Hurricanes	5	1622	24.31
Fires	5	3485	53.79