
Convolutional LSTMs for Cloud-Robust Segmentation of Remote Sensing Imagery

Marc Rußwurm

Chair of Remote Sensing Technology
Technical University of Munich
marc.russwurm@tum.de

Marco Körner

Chair of Remote Sensing Technology
Technical University of Munich
marco.koerner@tum.de

Abstract

Clouds frequently cover the Earth’s surface and pose an omnipresent challenge to optical Earth observation methods. The vast majority of remote sensing approaches either selectively choose single cloud-free observations or employ a pre-classification strategy to identify and mask cloudy pixels. We follow a different strategy and treat cloud coverage as noise that is inherent to the observed satellite data. In prior work, we directly employed a straightforward *convolutional long short-term memory* network for vegetation classification without explicit cloud filtering and achieved state-of-the-art classification accuracies. In this work, we investigate this cloud-robustness further by visualizing internal cell activations and performing an ablation experiment on datasets of different cloud coverage. In the visualizations of network states, we identified some cells in which modulation and input gates closed on cloudy pixels. This indicates that the network has internalized a cloud-filtering mechanism without being specifically trained on cloud labels. Overall, our results question the necessity of sophisticated pre-processing pipelines for multi-temporal deep learning approaches. We have published the source code of your work on GitHub¹.

1 Introduction

An increasing number of satellites monitor dynamic spatiotemporal processes on the Earth’s surface. Optical satellites measure spectral reflectances at multiple short electromagnetic wavelengths in regular time intervals of a few days. The surface, however, is often completely or partially covered by clouds. This poses an omnipresent challenge for the majority of remote sensing approaches designed with cloud-free imagery in mind. In prior published work [Rußwurm and Körner, 2018, 2017], we focused on the identification of crop classes with convolutional recurrent neural networks [Xingjian et al., 2015] to approximate life cycle events of vegetation. These phenological cycles and other land cover dynamics can be monitored at weekly intervals at spatial resolutions of several meters which allows distinguishing large single objects. We observed that state-of-the-art accuracies could be achieved without dedicated pre-processing mechanisms. In this work, we investigate this cloud-robustness further and treat clouds as data-inherent noise.

2 Related Work

Clouds distinguish themselves from ground pixels by their high reflectance. Rule-based models [Hollstein et al., 2016, Zhu and Woodcock, 2012, Hagolle et al., 2010] are used for many remote

¹<https://github.com/TUM-LMF/MTLCC>

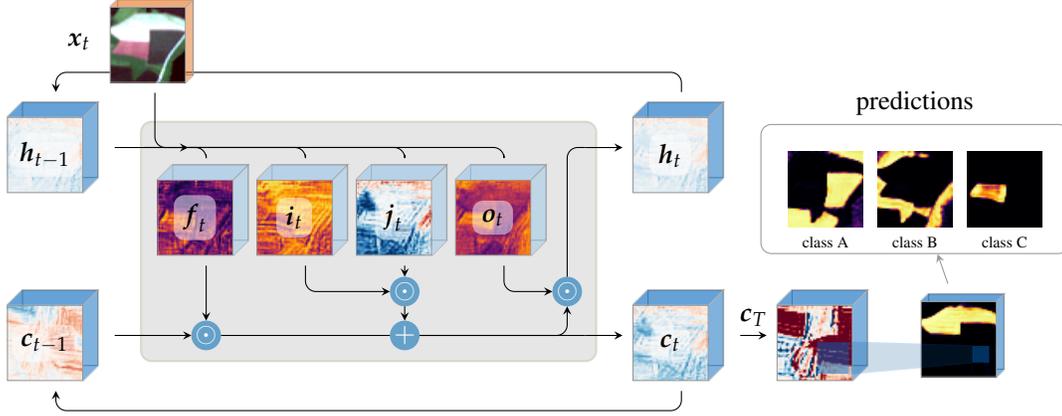


Figure 1: Illustration of the two-component convolutional long short-term memory network (LSTM) topology employed in this work. Each input image x_t of a series of T images is passed sequentially to the LSTM encoder. Classification relevant features are extracted to the internal cell state tensor c_T . A second convolutional layer compresses the dimensionality to the number of classes which yields activations per class.

sensing applications and rely on expert-designed features [Li et al., 2018]. The *fmask* algorithm [Zhu and Woodcock, 2012] and improved versions [Zhu et al., 2015, Frantz et al., 2015] implement a projection of the detected cloud on the surface to additionally predict the shadow cast by the cloud. Other approaches extract multi-temporal features and utilize the sudden temporal increase in reflectance [Hagolle et al., 2010]. *Convolutional neural networks (CNNs)* have also shown to compare well on cloud and cloud-shadow classification [Li et al., 2018] indicating that these features can be learned from the data by deep neural networks. However, masking single pixels by a pre-classification introduces an additional layer of complexity and raises the question of how to treat these masked pixels accordingly. For instance, one recently published approach [Interdonato et al., 2018] replaces cloudy pixels by temporal interpolation of cloud-free observations, while we introduced pre-classified cloud labels as additional prediction class in early work on this topic [Rußwurm and Körner, 2017]. Overall, cloud-filtering remains a pre-processing necessity for most remote sensing approaches that are prone to fail in the presence of data noise. Similar to our work, only a few approaches have tried to design robust methods that do not require this additional pre-classification step. For instance, Man et al. [2018] used ensemble-based methods of supervised classifiers which have shown a certain robustness to the presence of clouds.

3 Method

In this section, we outline the theoretical basis of *convolutional long short-term memory (convLSTM)* networks utilized in this work and provide detail on the employed network topology.

3.1 Convolutional Long Short-term Memory

Long short-term memory (LSTM) networks [Hochreiter and Schmidhuber, 1997] implement internal gates to control the gradient-flow through time and an additional container for long-term memory c_t . This yields the LSTM update $(h_t, c_{t-1}) \leftarrow (x_t, h_{t-1}, c_{t-1})$ that maps an input x_t and short-term context h_{t-1} to a hidden representation h_t . Additionally, a long-term cell state c_{t-1} is updated to c_t at each iteration and can store information for a theoretically unlimited number of iteration. Three gates control the update of the cell state $c_t \leftarrow c_{t-1} \odot f_t + i_t \odot j_t$ by element-wise multiplication \odot . The forget gate $f_t = \sigma(x_t * \theta_{fx} + h_{t-1} * \theta_{fh} + 1)$ evaluates the influence of the previous cell state c_{t-1} with a sigmoidal $\sigma(\cdot) \in]0, 1[$ activation function. The input and modulation gates

$$i_t = \sigma(x_t * \theta_{ix} + h_{t-1} * \theta_{ih}), \text{ and } j_t = \tanh(x_t * \theta_{jx} + h_{t-1} * \theta_{jh}) \quad (1)$$

are element-wise multiplied for the cell state update. The output gate $o_t = \tanh(x_t * \theta_{ox} + h_{t-1} * \theta_{oh})$ determines together with the cell state the current cell output

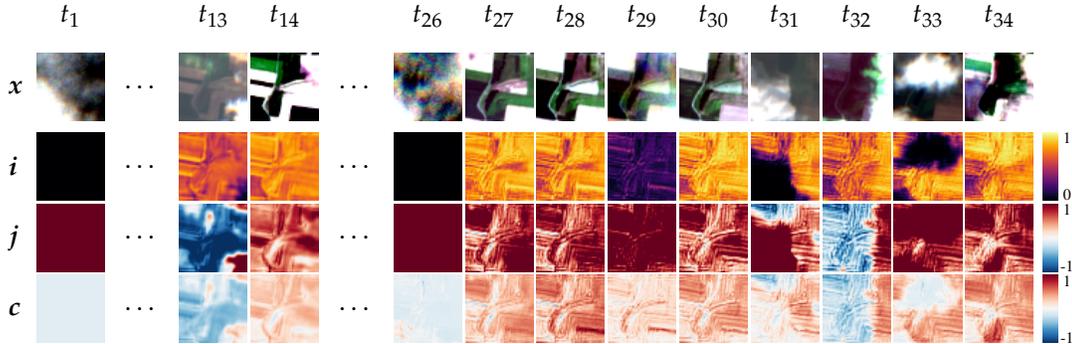


Figure 2: Activations of the cell state and selected gates of one convolutional LSTM cell that indicate that the cell has internalized a cloud-filtering scheme. The input gate i in this specific cell seems to be assigned values of zero on cloudy pixels as seen at steps $t = \{13, 26, 31, 33\}$.

$h_t \leftarrow o_t \odot c_t$. Convolutional recurrent networks implement a convolution, denoted by $*$, instead of a matrix multiplication. Each respective gate activation, referred by subscripts f, i, j, o , is controlled by trainable weights for input $\theta_{fx}, \theta_{ix}, \theta_{jx}, \theta_{ox} \in \mathbb{R}^{k \times k \times d \times r}$ and hidden representation $\theta_{fh}, \theta_{ih}, \theta_{jh}, \theta_{oh} \in \mathbb{R}^{k \times k \times r \times r}$ where d represents the number of image channels, k the convolutional kernel size, and r a hyper-parameter that determines the number of the hidden states in the recurrent layer. With this change, image data of certain width, height and depth can be processed with convolutions partially connecting the local pixel neighborhoods between layers.

3.2 Network Topology

We utilize this single-layer convolutional LSTM neural network to encode a sequence of T satellite images to the fixed length representation c_T , as illustrated in Fig. 1. In our initial work [Rußwurm and Körner, 2018], we found that 256 recurrent cells performed best and used this hyper-parameter for the number of hidden states within the ConvLSTM network. To balance the influence of the sequence order, we also encode the reversed sequence and append the final cell states. Applying softmax normalization produces activations that can be interpreted as network-confidences per class. We used convolutional kernels of $3 \times 3px$ throughout the network. To train, we evaluate the cross-entropy between the last layer and a one-hot representation of the ground truth labels. The influence of each weight on the evaluated loss is determined by back-propagated gradients which are used to iteratively adjust the respective weight using the Adam optimizer [Kingma and Ba, 2014] with a learning rate of $\lambda = 0.001$.

4 Results

The primary objective for this network was to identify the type of cultivated crops in an area of interest of $100 \text{ km} \times 40 \text{ km}$, as described at greater detail in Rußwurm and Körner [2018]. Hence, we trained our network end-to-end on label data describing the crop-type on distinct field parcels and provided no additional label information about cloud coverages. We used a sequence of 46 SENTINEL 2 satellite images from the year 2016 for this objective. This satellite measures the reflectances of 13 spectral bands at 10 m, 20 m, and 60 m resolution. To harmonize the data sources, we bi-linearly interpolated these to 10 m resolution and rasterized the crop labels accordingly. Additional information about area of interest and data partitioning strategy is given in Appendix A. In the remainder of this section, we evaluate the robustness of the proposed network to cloud coverage.

4.1 Exploration of Hidden LSTM States

We trained the network on field crop labels for thirty epochs using raw sequences of cloudy and non-cloudy observations. The top row of Fig. 2 shows an particular example sequence with $48 \times 48px$ images. The following rows illustrate activations of the internal convolutional LSTM for input gate

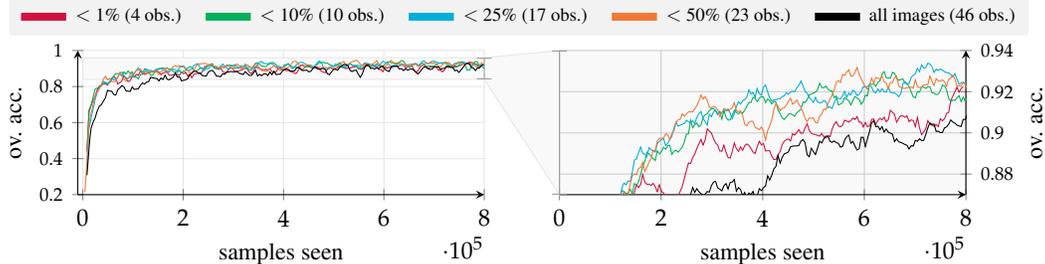


Figure 3: Overall accuracy over the training progress recorded on the validation partition with different degrees of cloud coverage. The network achieves similar accuracies on cloudy and non-cloudy data.

i , modulation gate j , and cell state c at each input time t . While all of the 256 recurrent cells likely contribute to the classification decision, only a few were visually interpretable similar to the shown example, as can be compared in Appendix B. The activations in Fig. 2 of these gates in the second and third row show that the input gate i approaches zero at pixels that are covered by clouds. This effect can be observed at time steps $t = \{13, 26, 31, 33\}$. At time $t = 32$ the input gate seems unchanged, however, the modulation gate j changes sign. Overall, these results indicate that the convolutional recurrent network has internalized a mechanism for cloud-filtering.

4.2 Ablation Experiment on Cloud Coverage

In this experiment, we trained the network on datasets with different degrees of cloud coverage. To determine the actual coverage of clouds per observation, all satellite images have been pre-processed using the `fmask` algorithm implemented in the `Sen2Cor` software, as is common practice in remote sensing [Conrad et al., 2014, Foerster et al., 2012]. With this, a cloud coverage pixel ratio per observation can be calculated. Based on this, several sub-datasets have been created with either all 46 observations, the 26 images covered with less than 50%, 17 images with less than 25%, 10 with less than 10%, and 4 completely cloud-free images. We trained the network on these pre-filtered datasets and show the overall accuracy recorded during training on validation data in Fig. 3. Even though the ratio of cloudy images varies greatly between these datasets, the classification accuracy remains similar for all of the sub-sampled datasets. The right graph shows a zoomed view and reveals some differences between the dataset performances. Datasets containing all observations and the four completely cloud-free observations have been slightly worse classified than the intermediate ones of 10%, 25%, and 50% coverage. It seems that the manual removal of the completely cloud-covered observations had a minor benefit on the classification accuracy. Also, the four cloud-free observations were classified slightly worse which indicates that these may have missed some characteristic vegetation-related events. Intuitively, these results show a trade-off between restrictions on cloud coverage and sequence length and demonstrate that cherry-picking single cloud-free observations may lead to inferior classification accuracy. Overall, these results demonstrate the robustness of the convolutional long short-term memory network to handle data containing temporal noise induced by cloud coverage.

5 Conclusion

Noise in temporal data is a common challenge for a variety of disciplines. In this work, we focused on noise induced by cloud coverage in multi-temporal remote sensing imagery. Most Earth observation approaches either select few completely cloud-free observations or use a pre-classification to mask cloudy pixels. The experiments of this work showed that this cloud-induced temporal noise can be suppressed by training on an other objective in an end-to-end fashion with an appropriate model design. Our results demonstrate that convolutional recurrent networks are able to consistently extract the classification-relevant features from observations between clouds without dedicated cloud labels. Overall, this work questions the necessity of sophisticated, partly hand-crafted pre-processing pipelines for remote sensing imagery.

References

- Christopher Conrad, Sebastian Fritsch, Julian Zeidler, Gerd Rücker, and Stefan Dech. Per-Field Irrigated Crop Classification in Arid Central Asia Using SPOT and ASTER Data. *Remote Sensing*, 2(4):1035–1056, apr 2010. ISSN 2072-4292. doi: 10.3390/rs2041035. URL <http://www.mdpi.com/2072-4292/2/4/1035/>.
- Christopher Conrad, Stefan Dech, Olena Dubovyk, Sebastian Fritsch, Doris Klein, Fabian Löw, Gunther Schorcht, and Julian Zeidler. Derivation of temporal windows for accurate crop discrimination in heterogeneous croplands of Uzbekistan using multitemporal RapidEye images. *Computers and Electronics in Agriculture*, 103:63–74, 2014. ISSN 01681699. doi: 10.1016/j.compag.2014.02.003. URL <http://dx.doi.org/10.1016/j.compag.2014.02.003>.
- Saskia Foerster, Klaus Kaden, Michael Foerster, and Sibylle Itzerott. Crop type mapping using spectral-temporal profiles and phenological information. *Computers and Electronics in Agriculture*, 89:30–40, 2012. ISSN 01681699. doi: 10.1016/j.compag.2012.07.015. URL <http://linkinghub.elsevier.com/retrieve/pii/S0168169912002013>.
- David Frantz, Achim Röder, Thomas Udelhoven, and Michael Schmidt. Enhancing the detectability of clouds and their shadows in multitemporal dryland landsat imagery: extending fmask. *IEEE Geoscience and Remote Sensing Letters*, 12(6):1242–1246, 2015.
- Olivier Hagolle, Mireille Huc, D Villa Pascual, and Gérard Dedieu. A multi-temporal method for cloud detection, applied to formosat-2, ven μ s, landsat and sentinel-2 images. *Remote Sensing of Environment*, 114(8):1747–1755, 2010.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- André Hollstein, Karl Segl, Luis Guanter, Maximilian Brell, and Marta Enesco. Ready-to-use methods for the detection of clouds, cirrus, snow, shadow, water and clear sky pixels in sentinel-2 msi images. *Remote Sensing*, 8(8):666, 2016.
- Roberto Interdonato, Dino Ienco, Raffaele Gaetano, and Kenji Ose. Duplo: A dual view point deep learning architecture for time series classification. *arXiv preprint arXiv:1809.07589*, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Zhiwei Li, Huanfeng Shen, Qing Cheng, Yuhao Liu, Shucheng You, and Zongyi He. Deep learning based cloud detection for remote sensing images by the fusion of multi-scale convolutional features. *arXiv preprint arXiv:1810.05801*, 2018.
- Chuc Duc Man, Thuy Thanh Nguyen, Hung Quang Bui, Kristofer Lasko, and Thanh Nhat Thi Nguyen. Improvement of land-cover classification over frequently cloud-covered areas using landsat 8 time-series composites and an ensemble of supervised classifiers. *International Journal of Remote Sensing*, 39(4):1243–1255, 2018.
- José M. Peña-Barragán, Moffatt K. Ngugi, Richard E. Plant, and Johan Six. Object-based crop identification using multiple vegetation indices, textural features and crop phenology. *Remote Sensing of Environment*, 115(6):1301–1316, 2011. ISSN 00344257. doi: 10.1016/j.rse.2011.01.009. URL <http://linkinghub.elsevier.com/retrieve/pii/S0034425711000290>.
- Marc Rußwurm and Marco Körner. Temporal vegetation modelling using long short-term memory networks for crop identification from medium-resolution multi-spectral satellite images. In *CVPR Workshops*, pages 1496–1504, 2017.
- Marc Rußwurm and Marco Körner. Multi-temporal land cover classification with sequential recurrent encoders. *ISPRS International Journal of Geo-Information*, 7(4):129, 2018.
- Shi Xingjian, Zhoung Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015.
- Zhe Zhu and Curtis E Woodcock. Object-based cloud and cloud shadow detection in landsat imagery. *Remote sensing of environment*, 118:83–94, 2012.
- Zhe Zhu, Shixiong Wang, and Curtis E Woodcock. Improvement and expansion of the fmask algorithm: Cloud, cloud shadow, and snow detection for landsats 4–7, 8, and sentinel 2 images. *Remote Sensing of Environment*, 159:269–277, 2015.

Appendices

This addendum provides further background information on the crop classification task out of which the presented evaluation on cloud-robustness has emerged.

A Dataset and Area of Interest

The area of interest (AOI) north of Munich, Germany, was divided into partitions for the training, validation, and evaluation, as shown in Fig. 4. This ensures independence between the optimization of network parameters, hyper-parameter tuning, and final evaluation. Further, we took special care to avoid introducing dependence between these partitions by spatial proximity by subdividing the AOI into rectangular blocks of $3.84 \text{ km} \times 3.84 \text{ km}$ (multiples of the chosen tile sizes 240 m and 480 m). Margins between blocks of 480m avoid that tiles of two distinct partitions are located in immediate proximity. These blocks were randomly assigned to the partitions in a 4:1:1 train/valid/eval ratio. All tiles (240m and 480m) that were fully within the respective blocks have been assigned to their respective partition. The overall cloud coverage over the area of interest is shown in Fig. 5.

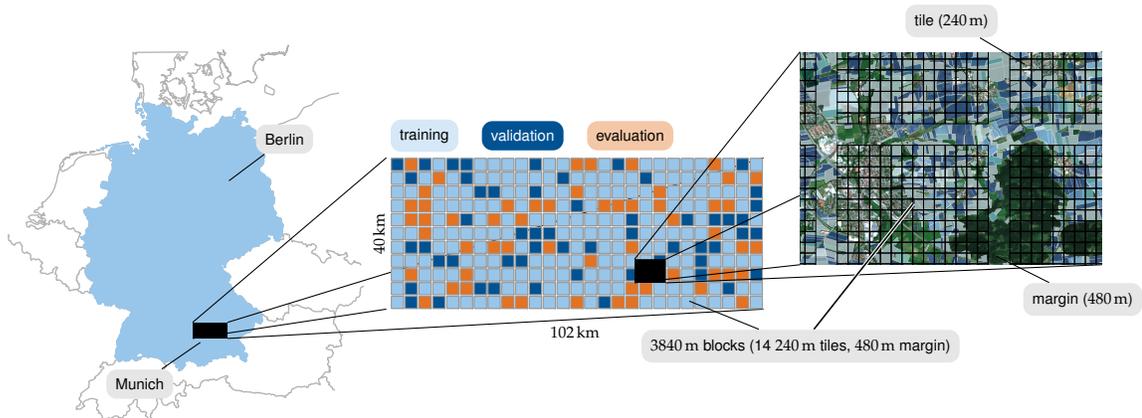


Figure 4: Overview over the area of interest (AOI)

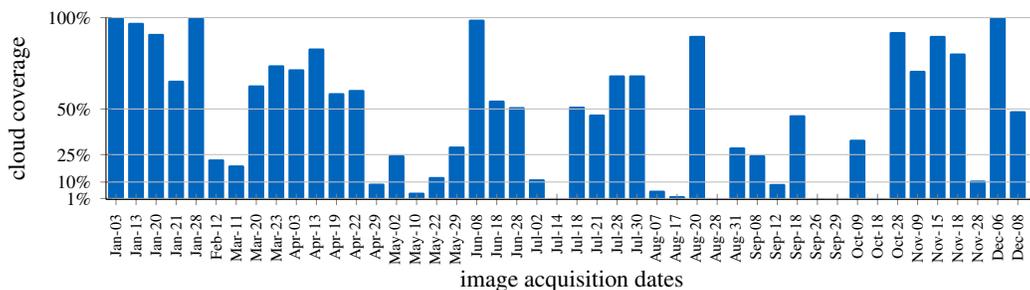
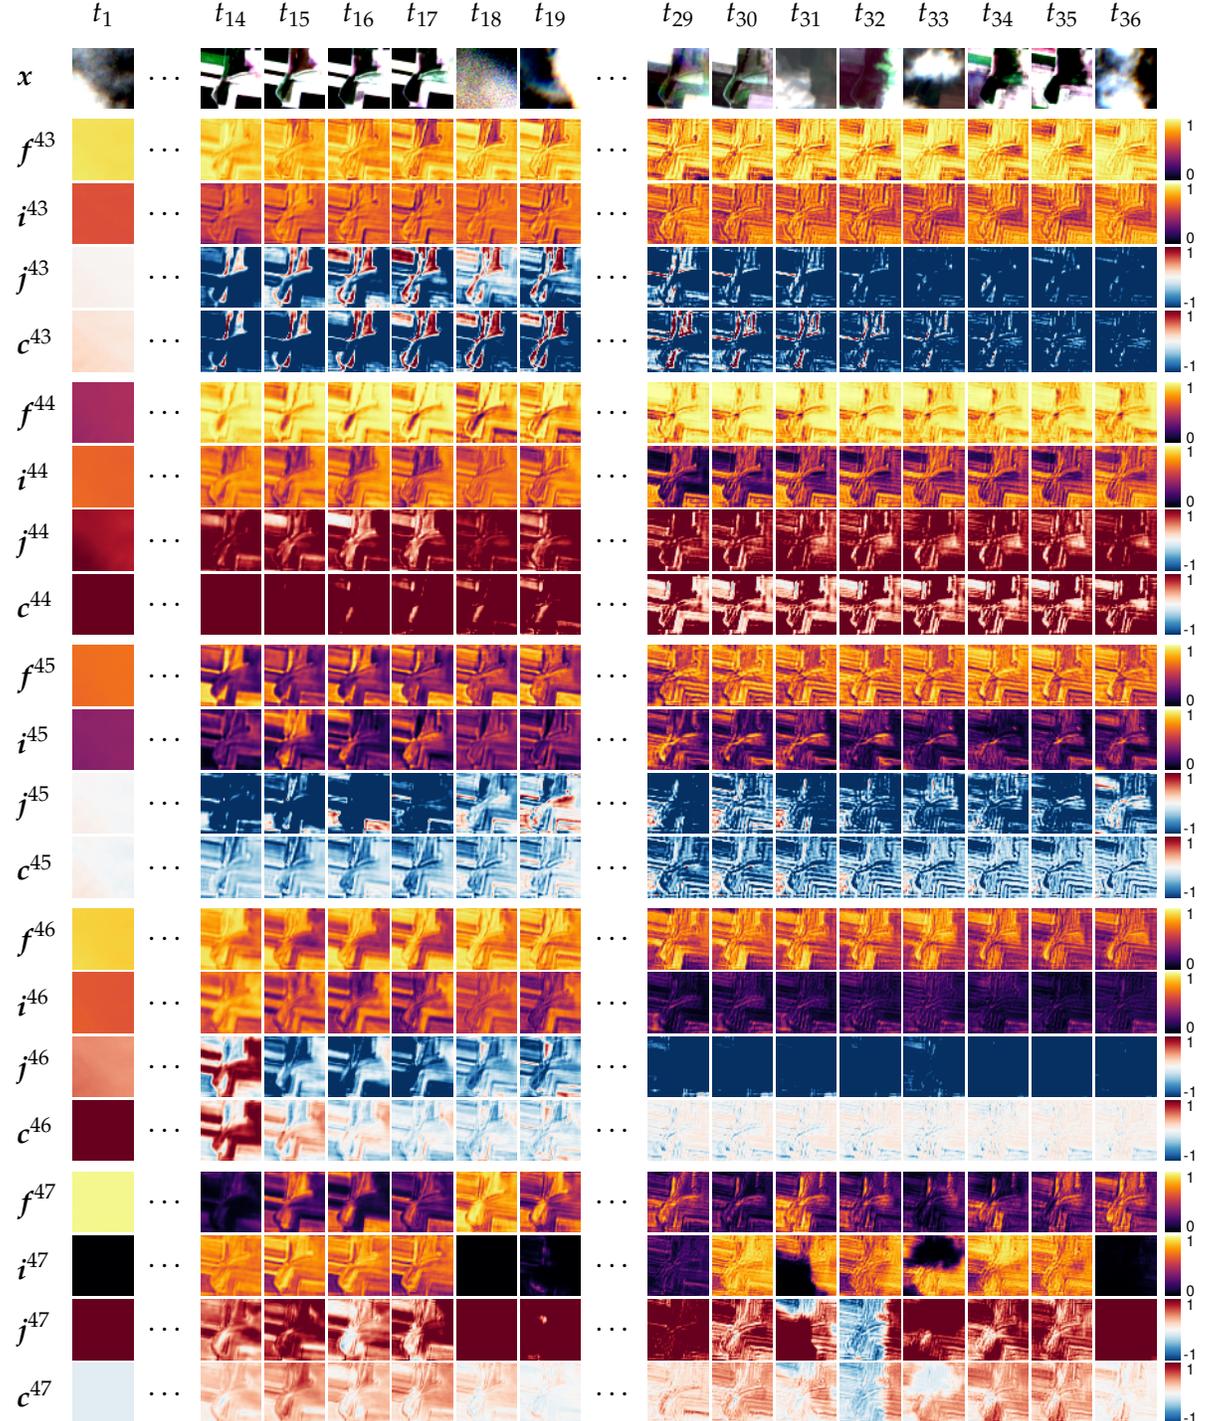


Figure 5: Actual cloud coverage over the area of interest obtained from the Sen2cor process for atmospheric correction.

B ConvLSTM Cell Activations

Further internal cell activation using the same network with 256 hidden states on two tiles. The top row shows the input image in an RGB representation. The gate activations of five hidden subsequent states are shown below. Hidden state 47 (denoted by raised index) shows a sensitivity to cloud coverage in both tiles.

B.1 Tile A



B.2 Tile B

