

TeleViT: Teleconnection-driven Transformers Improve Subseasonal to Seasonal Wildfire Forecasting

Ioannis Prapas^{1,3} Nikolaos Ioannis Bountos^{1,2} Spyros Kondylatos^{1,3}
Dimitrios Michail² Gustau Camps-Valls³ Ioannis Papoutsis¹

¹*Orion Lab, Institute for Astronomy, Astrophysics, Space Applications,
and Remote Sensing (IAASARS), National Observatory of Athens*

²*Department of Informatics and Telematics, Harokopio University of Athens*

³*Image Processing Laboratory (IPL), Universitat de València*

Abstract

Wildfires are increasingly exacerbated as a result of climate change, necessitating advanced proactive measures for effective mitigation. It is important to forecast wildfires weeks and months in advance to plan forest fuel management, resource procurement and allocation. To achieve such accurate long-term forecasts at a global scale, it is crucial to employ models that account for the Earth system's inherent spatio-temporal interactions, such as memory effects and teleconnections. We propose a teleconnection-driven vision transformer (TeleViT), capable of treating the Earth as one interconnected system, integrating fine-grained local-scale inputs with global-scale inputs, such as climate indices and coarse-grained global variables. Through comprehensive experimentation, we demonstrate the superiority of TeleViT in accurately predicting global burned area patterns for various forecasting windows, up to four months in advance. The gain is especially pronounced in larger forecasting windows, demonstrating the improved ability of deep learning models that exploit teleconnections to capture Earth system dynamics. Code available at github.com/Orion-Ai-Lab/TeleViT.

1. Introduction

Global warming increases the frequency and intensity of fire weather, amplifying the likelihood of the conditions that lead to extreme wildfire events [10]. In that context, it is important to improve our understanding of wildfires, anticipating wildfire patterns weeks and months in advance. Operationally, subseasonal to seasonal wildfire forecasts are typically treated as anticipated anomalies in temperature and precipitation, derived from process-based climate models [20], while disregarding crucial fire-related factors like soil moisture and vegetation dynamics. Deep Learning (DL) methods, able to learn from data, offer a promising avenue to

model Earth system processes such as wildfires [17], arising from the dynamic interactions of all the different fire drivers, namely climate, vegetation and human activity [8].

Previous work In fact, several studies have successfully applied DL to wildfire forecasting tasks [7,9]. When it comes to global predictions on subseasonal to seasonal scales, existing work relies on traditional Machine Learning (ML) approaches [16] that do not effectively capture Earth system dynamics that are important for long-term forecasting, even if sometimes they use teleconnection indices as input features [14,21]. In this work, we argue that for predictions on long temporal scales, it is crucial to train models that consider the Earth as one interconnected system, accounting for spatio-temporal interactions such as memory effects and teleconnections. There is substantial evidence that teleconnections modulate global wildfires [2,5,11,12]. For example, extreme wildfires in Siberia have been linked to preceding arctic oscillation [12] and previous-year soil moisture anomalies [5]. Despite the evidence, there is very limited ML work to predict wildfires that combines local information from the fire drivers with Earth system variables. Chen *et al.* [3] use simple autoregressive statistical methods to combine volume pressure deficit values with oceanic indicators. Yu *et al.* [21] use a statistical pre-processing to identify the most prominent oceanic indicators modulating burned area in Africa and then use the findings to select input features for tree-based machine learning models. AttentionFire [14] uses an attention-enhanced recurrent neural network architecture that considers temporal context and shows slight improvements at longer horizons when adding information from oceanic indices. More sophisticated models that handle the Earth as a system are arising naturally for weather and climate prediction. Earthformer [6] combines spatio-temporal attention with learnable global vectors that are meant to summarize the dynamics of a system. GraphCast [13] uses graph neural

networks that capture long-range interactions by defining graphs at several resolutions. Climax [15] is proposed as a transformer-based foundational model for weather and climate at coarse resolutions, working with global input or on a very large spatial domain.

Application context Our study has direct implications for humanitarian assistance and disaster response (HADR) operations, particularly in addressing the challenges of intensified wildfires due to climate change. Forecasting wildfires weeks or months in advance can inform the deployment of HADR-related anticipatory actions. This includes measures for vegetation and forest management (*e.g.* controlled burns and targeted firebreaks), coordination with local authorities to form evacuation plans, procurement of resources such as leasing firefighting equipment, and timely mobilisation of international and cross-border cooperation and aid.

Our transformer-based models combine local and non-local inputs like teleconnection indices and coarsened global variables, and thus offer a novel approach to model Earth system processes and improve our capabilities to forecast wildfires at subseasonal to seasonal scales. Outside of wildfire forecasting, these methods hold promise for diverse HADR applications, enabling foresight across extended timeframes, which can empower proactive planning and optimized resource allocation before disastrous events.

Contributions In this work, we propose to explicitly model short (local) and long (global) range interactions of the Earth system for long-term global wildfire pattern forecasting. For this, we develop a Teleconnection-driven Vision Transformer (TeleViT), that expands ViT [4] with an asymmetric tokenisation procedure in order to seamlessly combine local and global scale inputs, *i.e.* climate indices and coarsened global variables. We thoroughly examine the performance of the proposed model, demonstrating superior performance in various forecasting windows compared to models that do not leverage teleconnections.

2. Teleconnection Vision Transformer (TeleViT)

To produce reliable long-range forecasts (weeks or even months in advance), it is imperative to treat the Earth as a connected system. The Earth system is characterised by continuous interactions of processes that spread over large spatio-temporal windows, manifesting as memory effects and teleconnections. *Memory effects* are the realisation of the persistence or influence of past events on current and future states. They represent the temporal aspect of system behaviour, where the history of a system influences its present or future behaviour. In the context of wildfire prediction, memory effects capture how past events such as fuel accumulation, drought conditions, and weather patterns can impact

the likelihood, behaviour, and extent of wildfires. *Teleconnections* constitute long-distance interactions between different regions in the Earth system. They describe how changes in one region can influence atmospheric or oceanic conditions in another, often through large-scale atmospheric circulation patterns or oceanic phenomena [19]. Teleconnections are often described in terms of teleconnection indices, *i.e.* Oceanic and Climatic Indices (OCIs), calculated as large-scale anomalies of specific parameters, such as temperature, pressure or sea surface temperature. By teleconnections, however, we refer to the long-range spatio-temporal interactions and not to the indices per se, which we consider a mere proxy to the state of the Earth system.

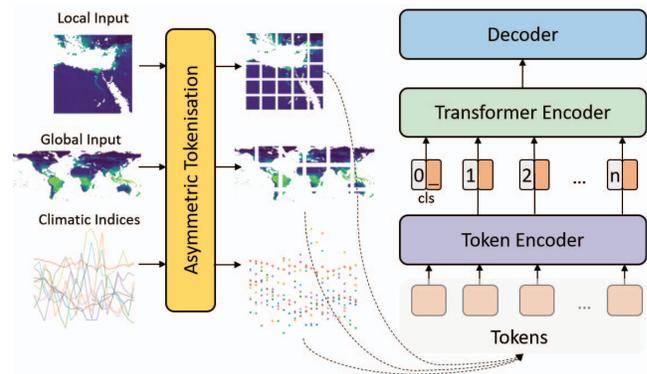


Figure 1. Full pipeline of the TeleViT architecture. The different multi-scale inputs *i.e.* local, global and teleconnection indices, are tokenized at different resolutions and fed to a Transformer encoder along with a prepended classification token. The linear decoder is based on the output of the classification token.

We propose to capture such distant interactions that are omnipresent in the Earth system, using data that can inform on the state and dynamics of the system. Particularly, we propose to enhance fine-resolution local data, *i.e.* information for a small area on Earth, that are commonly used in isolation, with information from i) OCIs and ii) coarse-resolution global data. Effectively combining these inputs is particularly challenging due to high dimensionality and the discrepancy between the spatial resolution of the different data sources.

In that direction, we propose to utilise the versatility of the Transformer architecture [18]. We build upon the Vision Transformer (ViT) [4], which adapts the architecture to computer vision problems by splitting an image into non-overlapping chips (tokens). The resulting token sequence serves as input to a standard Transformer. Extending ViT, we introduce the *Teleconnection-driven Vision Transformer (TeleViT)*. TeleViT relies on an asymmetric tokenization method which increments the sequence of input tokens with tokens stemming from different data sources, with potentially varying spatial and temporal resolutions. To produce the token sequence, each input source is tokenised inde-

pendently, taking into account its inherent characteristics along with the scale it operates on the Earth system. Local information and time series of climatic indices should be represented in high detail with smaller token sizes, while global information, operating at a greater scale, can be less detailed benefitting from larger token sizes.

TeleViT can be adapted to consume any type of input, given a tokenisation strategy. For our setting, we assume local input $x_l \in R^{C_l \times H_l \times W_l}$, global input $x_g \in R^{C_g \times H_g \times W_g}$ and indices input $x_i \in R^{C_i \times T}$, with C the number of features, H the height, W the width and T the time-series length. According to our asymmetric tokenisation, each of the inputs is split into tokens individually, resulting in N_l , N_g and N_i tokens with dimensions P_l , P_g and P_i that depend on the selected tokenisation strategy. Tokens are mapped to the embedding dimension D of the Transformer with a different trainable linear projection function for each input $f_k : R^{N_k \times P_k} \rightarrow R^{N_k \times D}$, $k \in l, g, i$. Similar to ViT, we add learnable positional encoding to the sequence of $N_l + N_g + N_i$ number of tokens, which is fed to a standard Transformer with depth K and A attention heads for each layer. Pixel-level prediction is performed by a simple decoder attached to the prepended classification token (cls). We choose a trainable linear layer that maps to the input resolution. The architecture is demonstrated in Figure 1.

As the Transformer architecture remains intact, most of the inductive biases come from the tokenisation procedure. By keeping the inductive bias to a minimum we remove any restrictions to known concepts, enabling unprecedented information combination of different input datasets, representing local and global inputs. For example, attention allows i) inter-dataset interactions, *i.e.* information flows from distant large-scale regions of coarse resolution to high-resolution localised windows, and ii) intra-dataset interactions, where several Earth system processes and interactions are modelled at a global scale while identifying location-specific characteristics for each region independently (depicted in Figure 4, Appendix D). The simplicity of the method makes it easily extensible to other data sources, as well as to the inclusion of a temporal dimension that we do not address in this study.

3. Experiments

We conduct our experiments on the SeasFire cube [1], a spatio-temporal dataset for subseasonal to seasonal wild-fire forecasting. It contains 21 years of data (2001-2021) at a global scale, in an 8-days temporal and 0.25° spatial resolution. The cube includes a diverse range of seasonal fire drivers, combining atmospheric, vegetation, and anthropogenic variables along with climate indices, in addition to target variables related to wildfires such as burned areas, fire radiative power, and wildfire-related emissions.

Building on the SeasFire cube, Prapas *et al.* [16] defined burned area pattern forecasting as a segmentation task. As

input, they use local patches that contain different channels of the fire driver variables and train a U-Net++ [22] to predict the presence of burned areas at a future time step. Their U-Net++ demonstrated a predictive skill greater than the burned area mean seasonal cycle for a lead forecasting time of up to 2 months. We follow a similar setup, where given a snapshot of the fire driver variables at timestep t we want to predict the presence of burned areas at a future timestep $t + h$. For the local input, we extract patches of size 80×80 , and as such, the world represented at 0.25° (1440×720 cells), is split into $18 \times 9 = 162$ local input patches. To extract the global input, we coarsen the cube to 1° (see Appendix C), reducing its size by a factor of 16, making the global input size 360×180 . Along with 10 fire driver variables extracted from the cube, the same for both local and global inputs, we calculate a *global positional encoding*, *i.e.* sine and cosine of the longitude and latitude. For each sample, 10 OCIs are extracted for the 10 months preceding t . The target is the presence of burned area at time-step $t + h$ for the region of the local input, where h is the lead time forecasting horizon. As such, a sample is comprised of four vectors; i) a local input x_l of size (14, 80, 80), ii) a global input x_g of size (14, 360, 180), iii) an OCI input x_i of size (10, 10) and iv) a target vector of size (1, 80, 80). The variables used are shown in Appendix A. The model’s performance is evaluated using the Area Under the Precision-Recall Curve (AUPRC). The train, validation, and test split is time-based, using years 2002 - 2017 for training, 2018 for validation and 2019 for testing. Only samples that contain burned areas are considered.

We assess the following models: i) a U-Net++, which uses only local input x_l [16], ii) a simple ViT, which uses only local input x_l , iii) TeleViT_{*i*}, which uses OCIs x_i along with local input x_l , iv) TeleViT_{*g*}, which uses only global input x_g along with local input x_l , and v) TeleViT_{*i,g*}, which uses both OCIs x_i and global input x_g along with the local input x_l . The hyper-parameters have been tuned for simple ViT and applied to TeleViT models that use the same core architecture. We provide detailed information on architectural and training choices in Appendix B.

The performance of the five models is examined in several forecasting horizons $h \in \{1, 2, 4, 8, 16\}$. A model that predicts at the maximum forecasting horizon, *i.e.* 16×8-days in advance, learns to predict the burned area pattern of a particular 8-day period approximately four months in advance. A different model is trained for each h , for a total of 25 experiments.

4. Results and Discussion

Figure 2 summarises the results of the experiments. In general, all the models exhibit a declining trend in performance as the forecasting window increases. The decline, however, is much less steep for teleconnection-driven models, than for the baselines (ViT and U-Net++). ViT proves

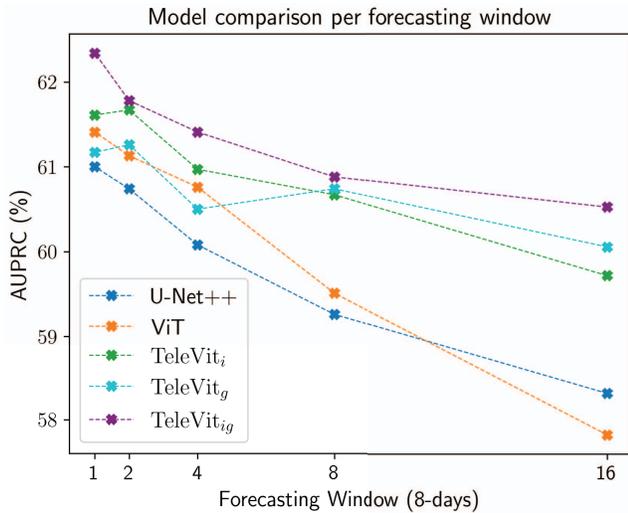


Figure 2. AUPRC performance of the different models for forecasting windows of 1, 2, 4, 8 and 16×8-days in advance.

to be a stronger baseline, outperforming U-Net++ with the exception of the 16×8-days, where U-Net++ overtakes it by a small margin. TeleViT_g achieves a comparable performance to ViT for short forecasting windows of up to 4×8-days, while it shows greater robustness to the increase of the forecasting window. TeleViT_i and TeleViT_{i,g} consistently surpass the baselines, with larger performance gaps as the lead time increases. Notably, TeleViT_{i,g} has the dominant performance, which indicates the benefits of a synergistic effect between teleconnection indices and global-scale representations. Interestingly, it achieves high gains even in short forecasting windows, which suggests that contextual information brought by this combination is helpful beyond teleconnections that operate on larger temporal scales. Figure 3 shows how the predictions compare to the target for a specific sample date, demonstrating high agreement.

In general, the results demonstrate TeleViT’s ability to successfully fuse local and non-local Earth system information and pave the road for exciting future research. Both global-scale inputs and OCIs can bring performance gains, but it is not clear if this is due to the same reasons. Further work is needed to investigate the relative contribution of each input. This urges us to comprehend the models, elucidating both familiar and undiscovered interactions between teleconnections and their influence on burned area patterns. A thorough examination of the attention maps may offer valuable insights in this regard, potentially shedding light on the underlying mechanisms. There is also much potential for further investigation of tokenisation schemes. The performance improvement induced by the introduction of coarsened global views may suggest that OCIs are a simplified proxy of the Earth system state that could be enhanced by global state representations. Further investigation will

reveal the contexts where this enhancement is most pronounced. Future work can exploit time series for both local and global inputs. In fact, it can be tested if global inputs with a temporal component can replace the indices. Finally, using models that treat the Earth as a system holds immense promise for scientific knowledge discovery, especially when incorporating knowledge of the physical systems.

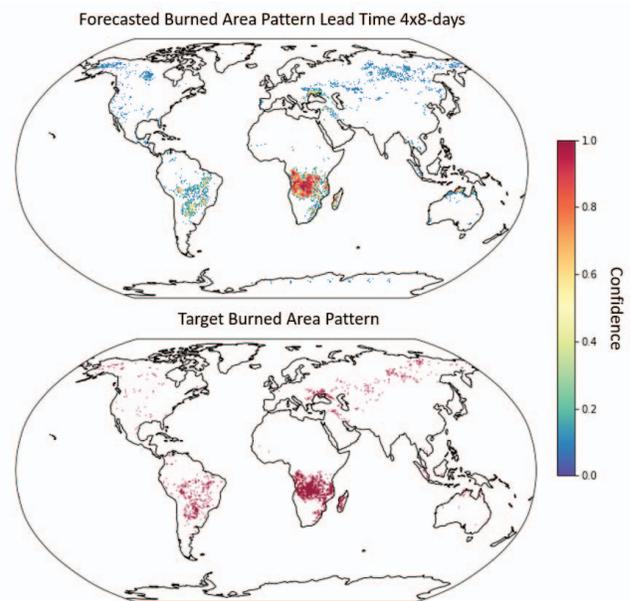


Figure 3. A sample prediction versus the target for the best model, predicting at 4×8-days lead forecasting time. Sea and values lower than 0.05 are masked out. Confidence is determined as the softmax score of the positive prediction.

5. Conclusions

In this work, we propose TeleViT, a transformer-based architecture that can leverage local and non-local information to model global Earth system processes. Using the SeasFire cube, we showcase our model’s ability to improve burned area pattern forecasting using climatic indices and coarsened global views of fire driver variables. This improvement is intensified for longer temporal forecasting horizons, where teleconnections are expected to have a more significant impact, strengthening our hypothesis that TeleViT can discover such long spatio-temporal Earth system interactions.

The combination of local information, climatic indices, and global views of the Earth holds significant potential for various HADR applications, promising enhanced anticipation capabilities at longer temporal horizons. This includes forecasting extreme events such as floods or droughts, the intensity of tropical cyclones, impacts on food security or human displacement. These applications can benefit from the enhanced prognosis, potentially improving proactive strategies and resource allocation in the face of disasters.

Acknowledgements

This work is part of the SeasFire project, which deals with "Earth System Deep Learning for Seasonal Fire Forecasting" and is funded by the European Space Agency (ESA) in the context of the ESA Future EO-1 Science for Society Call.

References

- [1] Lazaro Alonso, Fabian Gans, Ilektra Karasante, Akanksha Ahuja, Ioannis Prapas, Spyros Kondylatos, Ioannis Papoutsis, Eleanna Panagiotou, Dimitrios Mihail, Felix Cremer, Ulrich Weber, and Nuno Carvalhais. SeasFire cube: A global dataset for seasonal fire modeling in the earth system. **3, 6**
- [2] Adrián Cardil, Marcos Rodrigues, Mario Tapia, Renaud Barbero, Joaquin Ramírez, Cathelijne R. Stoof, Carlos Alberto Silva, Midhun Mohan, and Sergio de Miguel. Climate teleconnections modulate global burned area. *14(1):427*. Number: 1 Publisher: Nature Publishing Group. **1**
- [3] Yang Chen, James T. Randerson, Shane R. Coffield, Efi Foufoula-Georgiou, Padhraic Smyth, Casey A. Graff, Douglas C. Morton, Niels Andela, Guido R. Werf, Louis Giglio, and Lesley E. Ott. Forecasting global fire emissions on sub-seasonal to seasonal (s2s) time scales. *12(9)*. **1**
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Szekereit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. **2**
- [5] Matthias Forkel, Kirsten Thonicke, Christian Beer, Wolfgang Cramer, Sergey Bartalev, and Christiane Schmullius. Extreme fire events are related to previous-year surface moisture conditions in permafrost-underlain larch forests of siberia. *7(4):044021*. Publisher: IOP Publishing. **1**
- [6] Zhihan Gao, Xingjian Shi, Hao Wang, Yi Zhu, Yuyang Wang, Mu Li, and Dit-Yan Yeung. Earthformer: Exploring space-time transformers for earth system forecasting. **1**
- [7] Rafik Ghali and Moulay A. Akhloufi. Deep learning approaches for wildland fires using satellite remote sensing data: Detection, mapping, and prediction. *6(5):192*. Number: 5 Publisher: Multidisciplinary Digital Publishing Institute. **1**
- [8] Stijn Hantson, Almut Arneth, Sandy P. Harrison, Doug I. Kelley, I. Colin Prentice, Sam S. Rabin, Sally Archibald, Florent Mouillot, Steve R. Arnold, and Paulo Artaxo. The status and challenge of global fire modelling. *13(11):3359–3375*. Publisher: Copernicus Publications. **1**
- [9] Piyush Jain, Sean CP Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Steve Taylor, and Mike D. Flannigan. A review of machine learning applications in wildfire science and management. **1**
- [10] Matthew W. Jones, John T. Abatzoglou, Sander Veraverbeke, Niels Andela, Gitta Lasslop, Matthias Forkel, Adam J. P. Smith, Chantelle Burton, Richard A. Betts, Guido R. van der Werf, Stephen Sitch, Josep G. Canadell, Cristina Santín, Crystal Kolden, Stefan H. Doerr, and Corinne Le Quéré. Global and regional trends and drivers of fire under climate change. *60(3):e2020RG000726*. <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2020RG000726>. **1**
- [11] Flavio Justino, David H. Bromwich, Vanucia Schumacher, Alex daSilva, and Sheng-Hung Wang. Arctic oscillation and pacific-north american pattern dominated-modulation of fire danger and wildfire occurrence. *5(1):1–13*. Number: 1 Publisher: Nature Publishing Group. **1**
- [12] Jin-Soo Kim, Jong-Seong Kug, Su-Jong Jeong, Hotaek Park, and Gabriela Schaepman-Strub. Extensive fires in southeastern siberian permafrost linked to preceding arctic oscillation. *6(2):eaax3308*. Publisher: American Association for the Advancement of Science. **1**
- [13] Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Alexander Pritzel, Suman Ravuri, Timo Ewalds, Ferran Alet, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland, Jacklynn Stott, Oriol Vinyals, Shakir Mohamed, and Peter Battaglia. GraphCast: Learning skillful medium-range global weather forecasting. **1**
- [14] Fa Li, Qing Zhu, William J. Riley, Lei Zhao, Li Xu, Kunxiao-jia Yuan, Min Chen, Huayi Wu, Zhipeng Gui, Jianya Gong, and James T. Randerson. AttentionFire_v1.0: interpretable machine learning fire model for burned-area predictions over tropics. *16(3):869–884*. **1**
- [15] Tung Nguyen, Johannes Brandstetter, Ashish Kapoor, Jayesh K. Gupta, and Aditya Grover. ClimaX: A foundation model for weather and climate. **2**
- [16] Ioannis Prapas, Akanksha Ahuja, Spyros Kondylatos, Ilektra Karasante, Eleanna Panagiotou, Lazaro Alonso, Charalampos Davalas, Dimitrios Michail, Nuno Carvalhais, and Ioannis Papoutsis. Deep learning for global wildfire forecasting. **1, 3**
- [17] Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, and Nuno Carvalhais. Deep learning and process understanding for data-driven earth system science. *566(7743):195–204*. Publisher: Nature Publishing Group. **1**
- [18] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Szekereit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008. **2**
- [19] John M. Wallace and David S. Gutzler. Teleconnections in the geopotential height field during the northern hemisphere winter. *109(4):784–812*. Publisher: American Meteorological Society Section: Monthly Weather Review. **2**
- [20] Wikipedia contributors. Effis - seasonal forecast. <https://effis.jrc.ec.europa.eu/about-effis/technical-background/seasonal-forecast-explained>, 2023. [Online; accessed 16-June-2023]. **1**
- [21] Yan Yu, Jiafu Mao, Peter E. Thornton, Michael Notaro, Stan D. Wullschleger, Xiaoying Shi, Forrest M. Hoffman, and Yaoping Wang. Quantifying the drivers and predictability of seasonal changes in african fire. *11(1):2893*. Number: 1 Publisher: Nature Publishing Group. **1**
- [22] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. UNet++: Redesigning skip connections to exploit multiscale features in image segmentation. **3**

A. Input Data from the SeasFire cube

Table 1 shows which variables are used from the SeasFire cube [1], along with some pre-processing made for each variable. Total precipitation and population are log-transformed with $\log(1 + x)$ to follow a less skewed distribution. For more details on the variables, the reader is referred to the cited dataset.

Full name	Pre-processing
Local/Global Variables	
Mean sea level pressure	
Total precipitation	Log-transformed
Vapour Pressure Deficit	
Sea Surface Temperature	
Mean Temperature at 2 meters	
Surface solar radiation downwards	
Volumetric soil water level 1	
Land Surface Temperature at day	
Normalized Difference Vegetation Index	
Population density	Log-transformed
Cosine of longitude	Calculated
Sine of longitude	Calculated
Cosine of latitude	Calculated
Sine of latitude	Calculated
Climatic Indices	
Western Pacific Index	
Pacific North American Index	
North Atlantic Oscillation	
Southern Oscillation Index	
Global Mean Temperature	
Pacific Decadal Oscillation	
Eastern Asia/Western Russia	
East Pacific/North Pacific Oscillation	
Niño 3.4 Anomaly	
Bivariate ENSO Timeseries	
Target Variable	
Burned Areas from GWIS	Made binary

Table 1. Input and target variables used from the SeasFire cube for all settings. The same variables are used for both local and global views.

B. Model Details and Hyperparameters

Models are trained for 50 epochs. We use the cross-entropy loss and the Adam optimizer to train the models. For the U-Net++ model, the initial learning rate is set to 0.001, while for the Transformer models, it is set to 0.0001. The learning rate is reduced on the plateau and the weight decay is set to 0.000001. The model with the lowest validation

loss is used for testing. Before entering the models, local and global inputs are normalized. OCIs, which are anomalies are divided by their standard deviation.

The encoders of both ViT and TeleViT consist of $K = 8$ layers, with $A = 12$ attention heads each, and an embedding dimension $D = 768$. For the asymmetric tokenisation, we set $P_l = (1, 16, 16)$, $P_g = (1, 30, 30)$ and $P_i = 1$. This means the following:

- The local input x_l of size $(14, 80, 80)$ is tokenised spatially in 5×5 number of tokens with size 16×16 .
- The global input x_g of size $(14, 360, 180)$ is tokenised spatially in 12×6 number of tokens with size 30×30 .
- The OCI input x_i of size $(10, 10)$ is tokenised in 10×10 number of tokens.

C. Coarsening the SeasFire cube

The SeasFire cube is provided as a xarray-compatible file that is in Zarr format, which makes it easy to coarsen with xarray. All that is needed is to provide an aggregation function for each of the variables and a coarsening factor for each dimension. We use a coarsening factor of 4 along the longitude and latitude dimensions to convert the cube from 0.25° to 1° spatial resolution and all the input variables are mean-aggregated.

D. Interactions between the different input datasets

Figure 4 shows different interactions between the different input datasets captured by attention.

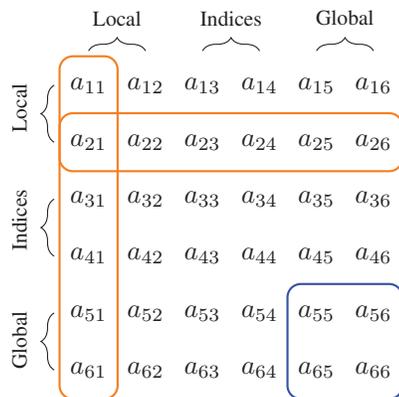


Figure 4. Depiction of intra-dataset (blue) and inter-dataset interactions (orange) in an attention matrix for a sequence of six tokens with an equal number of tokens for each data source.