
An Emulation Framework for Fire Front Spread

Andrew Bolt, Joel Janek Dabrowski, Carolyn Huston, Petra Kuhnert

Data61 / CSIRO, Australia

Abstract

Forecasting bushfire spread is an important element in fire prevention and response efforts. Empirical observations of bushfire spread can be used to estimate fire response under certain conditions. These observations form rate-of-spread models, which can be used to generate simulations. We use machine learning to drive the emulation approach for bushfires and show that emulation has the capacity to closely reproduce simulated fire-front data. We present a preliminary emulator approach with the capacity for fast emulation of complex simulations. Large numbers of predictions can then be generated as part of ensemble estimation techniques, which provide more robust and reliable forecasts of stochastic systems.

1 Introduction

Bushfires pose a serious threat to communities and natural flora and fauna across wide regions of Australia, as well as internationally. Simulated fires provide valuable data for first responders to assess vulnerable areas and the risk of a firefront impacting communities as well as being able to formulate response strategies for threatening fires. Simulation platforms such as Spark [Miller et al., 2015] and Phoenix [Tolhurst and Chong, 2010] use various techniques to predict how a fire front will progress through time. Underpinning such simulations are empirical rate-of-spread (ROS) calculations. These calculations determine how quickly a fire burns given a fuel source and varying environmental conditions such as temperature, slope, wind speed and wind direction [Cruz et al., 2015].

A single simulation instance generates one possible future fire front. To generate uncertainty estimates of a fire reaching a given location requires running ensembles of simulations under various environmental conditions. Generating large ensembles becomes computationally taxing and may be a prohibitive barrier to this type of analysis.

Emulation using machine learning is a method that attempts to mimic a slow running and highly parameterized process model using training examples. We develop an emulator that approximates a simulated fire front and discuss how surrogate models of this type could be used more efficiently in the future to characterise a broad range of fire scenarios simulated from varying environmental setups..

2 Modelling

2.1 Data

We use a data set of 200 simulated fires, generated using the Spark platform under real world meteorology and land input conditions in Australia. These trials are a subset of SPARK runs conducted as part of CSIRO's commercial work in the bushfire space. We aim to closely reproduce the simulated data using an emulator that requires a fraction of the computational resources. This makes running large ensembles to explore a broader range of fire scenarios, a more feasible proposition.

Input data for the simulation consists of a topographical map¹, weather data², and land classification map³. The resolution of spatial data is 30mx30m for each pixel. Weather station data is polled every 30 minutes.

Key pre-processing steps in our machine learning pipeline include: converting heightmaps to x and y gradient components using a Sobel edge algorithm; converting wind speed and direction to x and y components; expressing distances in pixel (30m) units and times in interval (30 minutes) units; and creating training samples by cropping simulated images to 256×256 pixel squares, centred about an active fire source. Cropping is used because the memory requirements for training neural networks using large images is prohibitive.

2.2 Neural Network Architecture

We approach the emulation of fire spread using a neural network (NN) framework due to their versatility as well as their success in structuring emulators for other problems [Sekou et al., 2019, Wang et al., 2019, Allaire et al., 2021, Burge et al., 2020]. The design of the NN should address some challenging aspects of the input feature and output space, such as image size, speed to generate outputs and a NN architecture that can cope with a mix of spatio-temporal inputs.

In terms of the size of the images that convey the fire front, each sample is allowed to vary, therefore accepting inputs to be of varying size. Fully convolutional networks are able to handle a variable input array size, so this is a natural choice.

A focus for emulation is speed and to ensure that the overall complexity of the network is minimized while still maintaining performance [Thiagarajan et al., 2020, Kasim et al., 2020]. Our approach is to downscale the spatial data (topography and land/fuel type) through strided convolution operators. This smaller set of latent features are then updated by each time interval until a final state is reached. This is far less computationally expensive than applying updates directly to a full sized array. The final latent state is then upscaled by transposed strided convolution operators where a final fire shape is output.

Finally, incorporating weather data is itself difficult since it more closely resembles time series data than image data. We could simply treat each data point as a uniform array of values and approach the problem using standard convolutional techniques. Undesirably, this greatly increases the number of convolutional operations that must be performed, which taxes memory and processing power. Instead, we transform the latent terrain layers so their depth dimension matches that of the weather input. By multiplication we transform these layers into an input with the correct spatial dimensions. This process is repeated for each new weather input until the final fire shape is produced.

Figure 1 shows a schematic for the neural network we deploy. A sample input image size of 256×256 pixels is used as a demonstration. The model uses a total of 106,532 trainable parameters. Note that there are only 21,248 trainable parameters in the residual block. The final model output layer is a spatial layer with values corresponding to arrival times of the fire.

2.3 Training the Neural Network

As the size of the images supplied for each fire simulation vary, some being quite large (up to 2048 pixels on an edge), we take the approach of cropping the images to 256 by 256 pixels in order to reduce the memory requirements needed for training. The cropping is centered around an actively burning region on the perimeter, and random rotation and flipping is performed.

The loss function, \mathcal{L} of an image, \mathcal{P} is represented as

$$\mathcal{L}(\mathcal{P}) = \log_{10} \left(\frac{MSE_o + \tau}{MSE_p + \tau} \right)$$

¹Topography data sets derived from Geoscience Australia SRTM-derived 1 Second Digital Elevation Models Version 1.0. Data is publicly available under Creative Commons Attribution 4.0 International Licence.

²Meteorological time series data sets derived from Bureau of Meteorology automated weather station data .

³Land classification data sets derived from Department of Agriculture and Water Resources (ABARES) Land Use of Australia 2010-11 data set. Data is publicly available under Creative Commons Attribution 3.0 Australia Licence.

where $MSE_o = \frac{1}{n} \sum_{i=1}^n (y_i^{(o)} - y_i)^2$, $MSE_p = \frac{1}{n} \sum_{i=1}^n (y_i^{(p)} - y_i)^2$ and $y_i^{(o)}$ is the initial observed fire arrival time at the i -th pixel, y_i is the target future arrival time at the i -th pixel, and $y_i^{(p)}$ is the predicted future fire arrival time at the i -th pixel. The term τ is a very small positive number to avoid asymptotic instability. The loss function is the log ratio of the mean squared error (MSE) of the initial and final arrival maps over the predicted and final arrival maps. This can be thought of as the improvement of the emulator over simply "doing nothing". The loss function was chosen since it does not weight samples with fast fire growth and more strongly than samples with limited fire growth.

The model is implemented using TensorFlow 2.0 and trained for 400 epochs. We used the ADAM optimizer [Diederik and Ba, 2015] and a batch size of 16. We withhold a test set with a split of 0.2. This set remains un-cropped.

3 Results

We present the model evaluation metrics in Table 3. There is a close agreement between results in the training and test sets, indicating that the model is generalising well. An additional benefit of cropping the training set is that it reduces the chance of over-fitting, and closely resembles principles from few-shot learning [Wang et al., 2021].

Figure 2 displays predictions for a test sample simulated over a duration of 30 minutes. In this sample there is overestimation of the fires spread, in which the emulator estimates the fire spreading faster than the simulation.

Figure 3 displays the fire front for this sample. Between the predicted and comparison (simulated) perimeters there is a good agreement of general shape. Broad features are in agreement, while intricate and narrow features present in the comparison perimeter are lost in the prediction. Importantly we see that the emulated behavior with respect to nonburnable terrain (darkest background shade) is consistent with the simulation. In particular we note that the fire does not advance North Easterly at the bend in the river on the left flank of the fire.

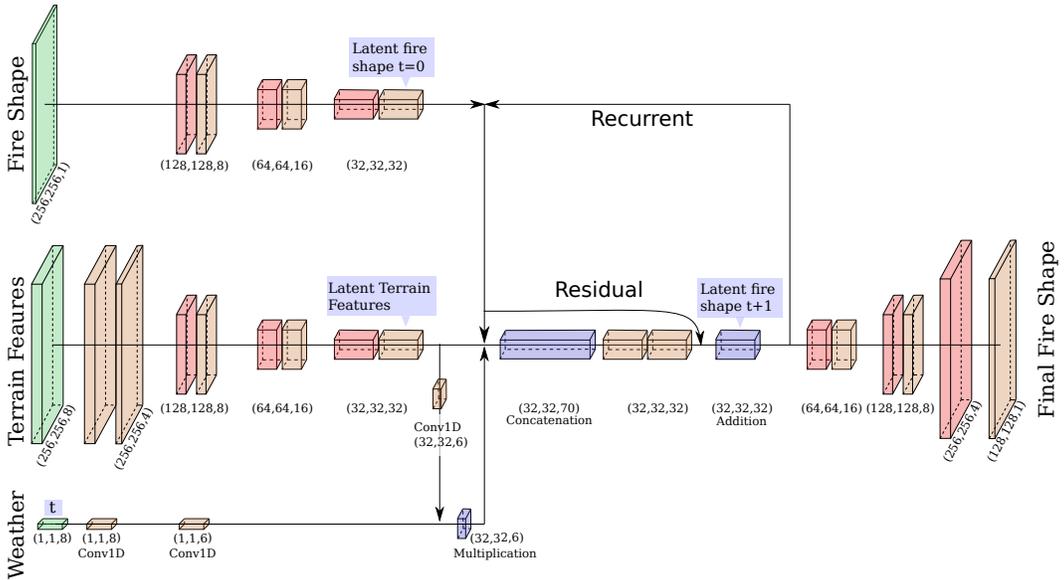


Figure 1: The emulator architecture. Red blocks represent strided convolutional (or transposed convolutional) layers with kernel size 4 and stride size 2. Orange blocks represent convolutional layers with kernel size 3 and stride size 1. The fire shape and terrain inputs are 2D arrays. In training we use 256x256 pixels, though the model is fully convolutional and can handle arbitrarily sized arrays. Weather inputs are a set of scalar values (eg. temperature, wind speed). These values are updated in successive steps as part of the input to the recurrent component of the model. The state of this recurrent cell are the latent fire shape features. Once the final weather input is processed upscaling is used to restore the final predicted fire shape.

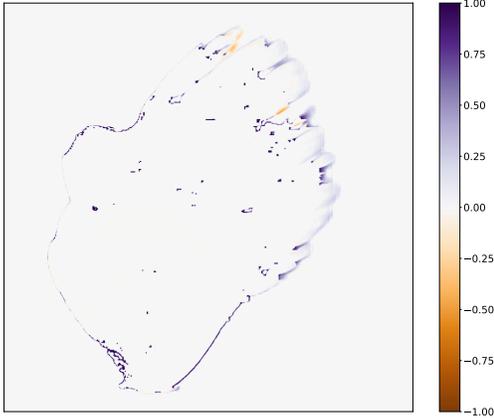


Figure 2: The difference between predicted and comparison (simulated) fire arrival times for a single interval (30 minutes) on a test sample. The interval is normalised so 0 represents the beginning and 1 the final arrival time. Positive values (purple) represent false-positives and negative values (orange) represent false-negatives.

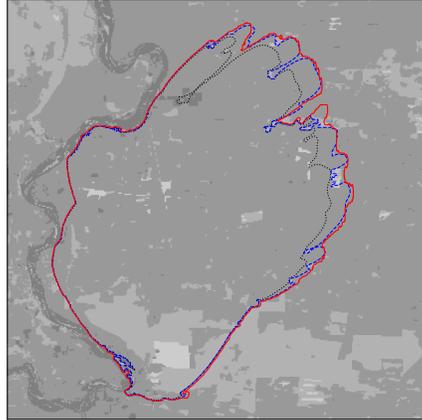


Figure 3: The initial fire front (dotted, black) shown against the predicted fire perimeter (solid, red) and the comparison fire perimeter (dashed, blue) for the same sample in Figure 2. These perimeters are overlaid on the various land classes used by the model, shown in gray-scale.

Set	Loss	Jaccard Score	Dice Score
training	-0.45	0.68	0.81
test	-0.49	0.67	0.79

Table 1: Model loss and evaluation metrics. Aggregate over all samples.

The recurrent component of our neural network allows for a series of intervals to build out a longer duration prediction. This is likely to be a better gauge of the emulator’s performance and is an extension that we are currently exploring.

In terms of bench-marking speed and memory requirements against conventional simulations we are still awaiting a more sophisticated analysis. Preliminary trials show that there is a speedup of around a factor of four. This may improve as the model architecture is refined, and the implementation is improved.

4 Conclusion

In this paper we have shown how a neural net can be constructed to efficiently emulate a spatio-temporal spread model. In this case we specifically focus on the emulation of fire front spread. While this work represents a preliminary investigation we show a respectable match between emulated and simulated results. If development of the emulator leads to a much faster representation of the physical process then this opens up a number of possible applications.

Of immediate interest is the use of emulators in ensemble forecasting and the generation of confidence intervals for fire front predictions. This approach allows the estimation of the likelihood of a fire reaching an area, rather than simply calculating the most likely fire front. Another area of interest is transfer learning to fine tune the model with the use of real fire examples. This could in principle lead to a neural network model that is more accurate than the original model and simulations that the emulator was developed on.

Broader Impact

In this paper we propose an architecture for emulating fire simulations from SPARK. The downscaling layers act to compress the data. This allows for a speedup over simulated fires which act on the full

uncompressed topographical scale. The reduced memory requirements may also be useful since large scale SPARK simulations are often constrained by memory.

We have shown a modest reduction in processing time between our emulations and SPARK simulations. It is likely that these gains will improve as the code and architecture are further refined.

While we have demonstrated that emulation shows promise as a methodology for mimicking fire spread, this geo-spatial approach may have impact in a variety of similar problem spaces. Areas such as disease spread, pollutant spread, and pest spread all represent similar problem scopes where this emulation approach may be viable.

References

- Claire Miller, James Hilton, Andrew Sullivan, and Mahesh Prakash. SPARK – a bushfire spread prediction tool. In *International Symposium on Environmental Software Systems*, pages 262–271, 2015. DOI: [10.1007/978-3-319-15994-2_26](https://doi.org/10.1007/978-3-319-15994-2_26).
- Kevin Tolhurst and Derek Chong. *PHOENIX RapidFire User Manual*. University of Melbourne, 2010.
- Miguel G Cruz, James Stanley Gould, Martin E Alexander, Andrew L Sullivan, W Lachlan McCaw, and Stuart Matthews. *A guide to rate of fire spread models for Australian vegetation*. CSIRO and AFAC, 2015.
- Taibou Sekou, Moncef Hidane, Julien Olivier, and Hubert Cardot. From patch to image segmentation using fully convolutional networks - application to retinal images. *arXiv:1904.03892*, 2019.
- Chang Wang, Zongya Zhao, Qiongqiong Ren, Yongtao Xu, and Ti Yu. Dense U-net based on patch-based learning for retinal vessel segmentation. *Entropy*, 21(2), 2019. DOI: [10.3390/e21020168](https://doi.org/10.3390/e21020168).
- Frédéric Allaire, Vivien Mallet, and Jean-Baptiste Filippi. Emulation of wildland fire spread simulation using deep learning. *Neural Networks*, 141:184–198, 2021. DOI: [10.1016/j.neunet.2021.04.006](https://doi.org/10.1016/j.neunet.2021.04.006).
- J Burge, M Bonanni, M Ihme, and Lily Hu. Convolutional LSTM neural networks for modeling wildland fire dynamics. *arXiv:2012.06679*, 2020.
- Jayaraman Thiagarajan, Bindya Venkatesh, Rushil Anirudh, Peer-Timo Bremer, Jim Gaffney, Gemma Anderson, and Brian Spears. Designing accurate emulators for scientific processes using calibration-driven deep models. *Nature Communications*, 11, 2020. DOI: [10.1038/s41467-020-19448-8](https://doi.org/10.1038/s41467-020-19448-8).
- M Kasim, D Watson-Parris, L Deaconn, S Oliver, F Hatford, D. H Froula, G Gregori, M Jarvis, S Khatiwala, J Korenaga, J Topp-Mugglestone, E Viezzer, and S. M Vinko. Building high accuracy emulators for scientific simulations with deep neural architecture search. *arXiv:2001.08055v2*, 2020.
- Kingma Diederik and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference for Learning Representations (ICLR)*, 2015.
- Yaqing Wang, Quanming Yao, James T. Kwok, and Lionel M. Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM Computing Surveys*, 53(3):1–34, 2021. DOI: [10.1145/3386252](https://doi.org/10.1145/3386252).