Deep Learning for High-Resolution Wildfire Modeling

Mark A. Finney[1*] , Jason M. Forthofer[1], Xinle Liu[2], John Burge[2], Matthias Ihme[2,3], Fei Sha[2], Yi-Fan Chen[2], Jason Hickey[2], John Anderson[2]

[1]*USDA Forest Service, Missoula Fire Sciences Laboratory, 5775 Highway 10 West, Missoula MT 59808*
{ mark.finney@usda.gov, jason.forthofer@usda.gov }
[2]*Google LLC, 1600 Amphitheatre Pkwy, Mountain View, CA 94043, USA, {liuxl@google.com, lawnguy@google.com, matthiasi@google.com, fsha@google.com, yifanchen@google.com, jyh@google.com, janders@google.com },*
[3]*Stanford University, 440 Escondido Mall, Stanford, CA 94305, USA*

**Corresponding author*

Keywords

Deep learning, wildfire spread modelling, fire behavior

Abstract

We demonstrate the use of a deep learning (DL) approach for representing the behavior of a high-resolution physics-based wildland fire spread model. The ultimate objective is being able to efficiently use the DL model for intensive simulations of large fires while retaining fidelity to the fine-scale physical processes. We begin with a fire model that reduces the spatial domain of the fire spread problem to one dimension (1D). The 1D model explicitly resolves cm-scale fuel variations, heat transfer and heating/drying dynamics of individual fuel particles and burning behavior of the bed. We then ran the fire model for 78,125 factorial combinations of fuel, weather, and topographic conditions as training data for the DL algorithm. The results of the DL analysis show overall agreement of 96% of the variation in fire behavior as represented by steady state rate of spread, flame length and flame zone depth. Exceptions to the DL regression indicate areas where more work is required in refining the resolution in training cases and use of advanced methods of embedding the fire model inside the DL algorithm loop.

1. Introduction

A major challenge for wildfire modelling is the ability to span the range of scales of fire phenomena (Clark et al. 2003). Fuel particle heat transfer and ignition occur over millimetres but wildfires impact landscapes and communities across kilometres. Even if the fine-scale physical processes of fire spread were well known, it is impractical to resolve them computationally for large domains of real fires. Thus, physics-based modelling has been much slower than real time and has not yet become operational for wildfire management. Even for research purposes on relatively small fires (~10^3 m^2) using supercomputing clusters, physical models employing computational fluid dynamics in 3D must compromise spatial resolution to achieve performance and thus fidelity to the fine-scale fuel descriptions, and processes of heat transfer, ignition, and combustion.

An alternative approach, long used in fire spread modelling, is to limit the spatial domain to one-dimension (1D) (Fons 1946). A 1D formulation resolves the heat transfer and sequential ignition of discrete fuel particles at fine-scales and can be unambiguously compared with measured behaviours. In a reprise of this 1D approach, Finney et al. (2021) describe the new research required to understand fuel burning rates, convection heat transfer, and heating and ignition of fuel particles at the fine scales. The fine scale explicitly represents fuel heterogeneity and employs a 1D model for heating, drying, and pyrolysis of fuel particles. It solves a single 1D line spanning domains of about 50-100 m in less than a minute, but this is too slow to be implemented directly in simulations of large fires.

Here we describe the application of deep learning (DL) techniques to the problem of representing modelled fire behaviour. Assuming our high-resolution 1D fire model captures the system behaviour accurately, DL can be used to determine the relations between the resulting fire behaviour outputs and a large list of input variables. Rapid advances in the past decade in DL show it can represent complex and perhaps unknown relations among variables in the system.

A rigorous background on DL is outside the scope of this paper. We refer the interested reader to (Goodfellow et al. 2016). Briefly, DL is a branch of machine learning (ML) that relies on using neural networks with multiple *hidden layers* of nodes that do not directly correspond to actual observations. These hidden layers allow DL models to build up rich latent representations of complex systems and are often distinguished from more traditional ML approaches in that they require less feature engineering to be successful, making them ideal candidates for complex domains in the geosciences.

There are many examples of traditional ML applications in wildfire modelling, but more modern DL applications are rare, particularly those modelling fire spread and growth (Piyush et al. 2020). Some examples include Hodges and Lattimer (2019) where a DL model was used to predict the structural evolution of wildfires in two dimensional simulations. Radke et al. (2019) used a DL model to predict how a single real-world fire evolves over the course of a 24-hour period. In a similar vein, we demonstrate that DL can be used to effectively model steady state variables predicted by the 1D fire spread model described below.

2. Methods

2.1. One Dimensional Fire Spread Model

The dynamical nonlinear 1D fire spread model is explained in detail by Finney et al. (2021) and in a companion paper that overviews a more recent version (Forthofer et al*., this conference*). Briefly, the model domain is a 1D transect through a fuel bed with resolution nominally of 1 cm. At these resolutions, variations in fuel structure, such as gaps and clumps can be explicitly resolved in contrast to bulk fuel descriptions (i.e., Scott and Burgan 2005). The simulation proceeds from a specified ignition at one end of the domain, calculating pyrolysis production and flame dimensions from fuel particles heated by radiation (flame and solid fuel) and convection. Convection heat flux to fuel particles from flames is calculated using a mean gas temperature profile and flame tip velocity (Finney et al. 2021). Fuel particle heating and drying response to steep heat flux gradients of an approaching fire (Cohen and Finney 2022a, b) are captured by a 1D conduction model for each size and material in each cell. Fire spread is an outcome of discrete particle ignition, which accelerates from the ignition state and achieves a steady rate of spread (ROS), flame length (FL), and flame zone depth (FZD).

To generate a dataset for training a DL model, the 1D fire spread model was run on a factorial combination of five values for each of the seven input variables, resulting in 78,125 distinct cases, see Table 1 and Table 2 for more details. For this demonstration, we chose only to represent steady state variables of fire spread rate, flame length, and flame zone depth from those data produced by the dynamical model. Others could include acceleration time and energy release from post-frontal combustion.

2.1. Pre-processing

The dataset is filtered to only keep data points where the 1D model reported realistic ROS values of no more than 10 m/s. The dataset is further normalized before model training such that all input variables are normalized into [0, 1] linearly between 0 and the max value of corresponding variable, and output variables are first square-rooted and then normalized in the same fashion as they're skewed and have a long tail. The dataset is randomly split into train (65%), validation (15%) and test datasets (20%) for model development before any filtering. The train dataset is used to learn the parameters of the DL model, while the validation dataset is used to select the best model to avoid overfitting the train dataset. The test dataset is never seen during model development and is reserved to measure the generalization quality of the DL model in the end.

Table 1. Descriptive statistics of input and output variables of the 78,125 full factorial combinations for training the DL model.

Category	Variable	Values (Full factorial design)	Train mean (std)	Validation mean (std)	Test mean (std)
Fuel	Particle Diameter (m)	{0.001, 0.002, 0.004, 0.005, 0.006}	0.0035 (0.0018)	0.0036 (0.0018)	0.0035 (0.0018)
	Moisture (%)	{5, 10, 15, 20, 25}	14.9 (7.0)	15.1 (7.0)	15 (7.0)
	Loading (kg/m3)	{0.1, 0.825, 1.55, 2.275, 3}	1.54 (1.02)	1.54 (1.02)	1.55 (1.02)
	Fuel Depth (m)	{0.01, 0.133, 0.255, 0.378, 0.5}	0.255 (0.173)	0.255 (0.173)	0.255 (0.172)
Weather	Wind Speed (m/s)	{0, 1.25, 2.5, 3.75, 5}	2.49 (1.76)	2.5 (1.77)	2.5 (1.77)
Topography	Terrain Slope (degree)	{0, 7.5, 15, 22.5, 30}	14.9 (10.5)	15.1 (10.6)	15 (10.5)
Fire	Bed Width (m)	{10, 20, 30, 40, 50}	30 (14.1)	29.7 (14.0)	29.9 (14.1)
Dependent variables	ROS (m/s)	N.A.	1.95 (10.11)	1.87 (8.29)	1.93 (10.38)
	Flame Length (m)		3.96 (2.57)	3.95 (2.59)	3.95 (2.56)
	Flame Zone Depth (m)		12.6 (12.0)	12.6 (12.1)	12.5 (12.0)

Table 2: Number of samples for the whole dataset (all) or each split (train, validation, and test). Each row corresponds to a different filtering, no filtering, valid ROSs or zero ROSs only.

Samples	All	Train	Validation	Test
All	78,125	50,681	11,712	15,732
Valid ROS	68,750	44,606	10,302	13,842
ROS == 0	6,666	4,307	1,020	1,339

2.2. Deep Learning Model

We use a two staged approach for predicting the steady-state variables. First, we use a binary classification model that predicts whether there is any fire at all (meaning the fire does not spread under the given conditions). If there is a prediction of *no fire*, then the ROS, FL and FZD are all zero. Otherwise, the regression model is used to perform the final prediction.

Given the relatively few inputs to the model, we chose the quintessential Feed-Forward Neural Network (Goodfellow et al. 2016, Chapter 6) (FFNNs) for both models. More specifically, we use a fully connected FFNN, which allows for a rich latent space that grows exponentially with the number of input features. As such, it is not appropriate for domains with large numbers of input features, but often works surprisingly well for smaller domains (Figure 1).

The DL models are trained on the training dataset such that the Mean Squared Error (MSE) (or binary cross-entropy, i.e., logistic loss) for the regression (or classification) model is minimized. L2 regularization was applied to avoid model overfitting. The learning rate and the learning rate decay rate were empirically determined by training multiple models and then selecting the models with the lowest loss on the *validation* dataset.

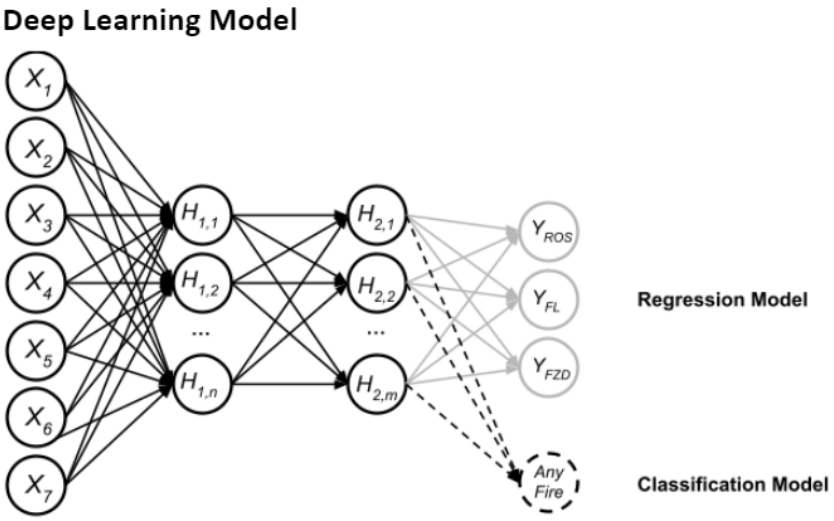

Figure 1. Architecture of feed-forward neural network used for classification and regression, consisting of two hidden layers. Common structures across both models are shown in solid black. The structure unique to the regression model is shown in solid gray. The structure unique to the classification model is shown in dashed-black. Both models take the same seven variables as input that feed into two hidden layers. The classification model requires fewer hidden nodes (n=8, m=8) whereas the regression model benefited from more nodes (n=32, m=8), and they had 145 and 547 parameters respectively. All nodes in the graph are perceptrons (the standard neural network unit) except for the Any Fire node, which also applies a final sigmoid transformation to facilitate binary classification.

3. Results

Use of the detailed 1D fire model as part of a large-scale wildfire simulation (2D or 3D) would provide the correct physics on combustion and steady state statistics but is infeasible due to the computational cost. Therefore, the primary goal for the DL model is to approximate the 1D model as accurately as possible at a dramatically reduced computational cost. For comparison, the average 1D model run time for the training data set was about 70 s for each of the 78,125 cases, whereas the DL model takes about 0.01 s for each case.

We show that the DL model captures the wildfire dynamics well (Figure 2). The scatter plots of the predicted dependent variables versus those from the underlying physical model for the test dataset show that most of the data points are quite close to the perfect reference orange (y = x) line. The mean (standard deviation) of the absolute errors for the dependent variables of ROS, FL, and FZD are 0.078 (0.17) m/s, 0.21 (0.29) m and 0.98 (1.53) m, respectively. Those numbers are close to that for the validation dataset, demonstrating that the DL model is capable of accurately reproducing the fire model's behaviour.

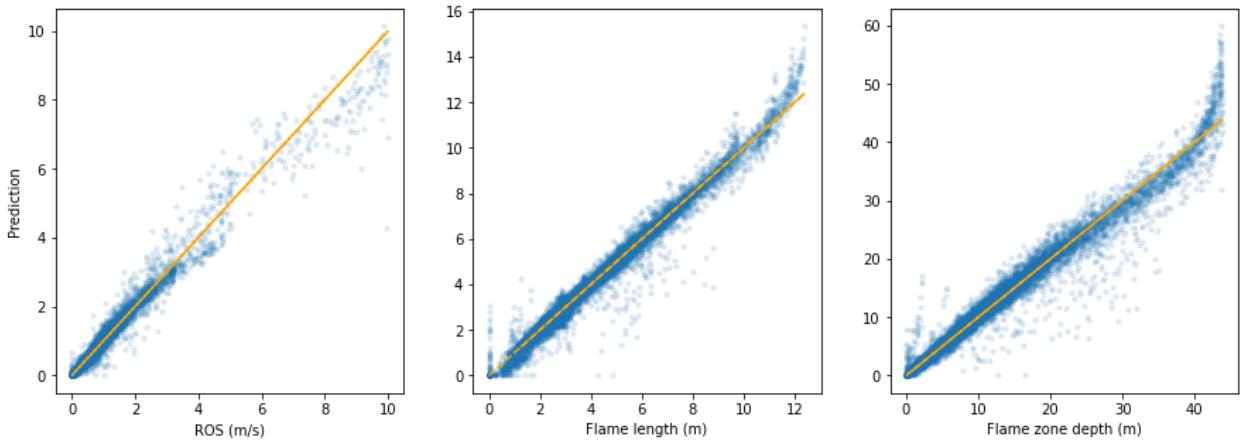

Figure 2: Scatter plot of predictions versus the values from the high-resolution physical model, for all three steady state variables, rate of spread, flame length and flame zone depth.

3.1. Qualitative Analysis

The example fire characteristics represented by ROS, FL and FZD are key descriptors of frontal fire behaviour. We found that errors made by the DL model are at the extremes of the input data. The error around zeros is an indicator that more work is required for the classification model in differentiating the no-fire cases correctly, which is an important problem in practical fire management. The overall underestimate for large ROSs and overestimate for large flames (Figure 2) might be related to dataset filtering, as the wind speed is at most 5 m/s, while we keep ROSs up to 10 m/s for cases with combined wind and high slope angle.

The DL model captured the monotonic dependence of fire behavior on the wind speed and terrain slope (Figure 3). It also and illustrates some extrapolation beyond the training data (recall that the whole dataset training data included only wind speeds up to 5 m/s), not only at the discrete wind speeds reported by the 1D model, but also over the continuous range. However, inevitably, at some point the models' limitations become apparent. Extending the models' accuracy to higher wind speeds can potentially be resolved by including training data at higher wind speeds or considering a probabilistic DL model, e.g. (Dillon et al. 2017), both of which are subjects of ongoing work.

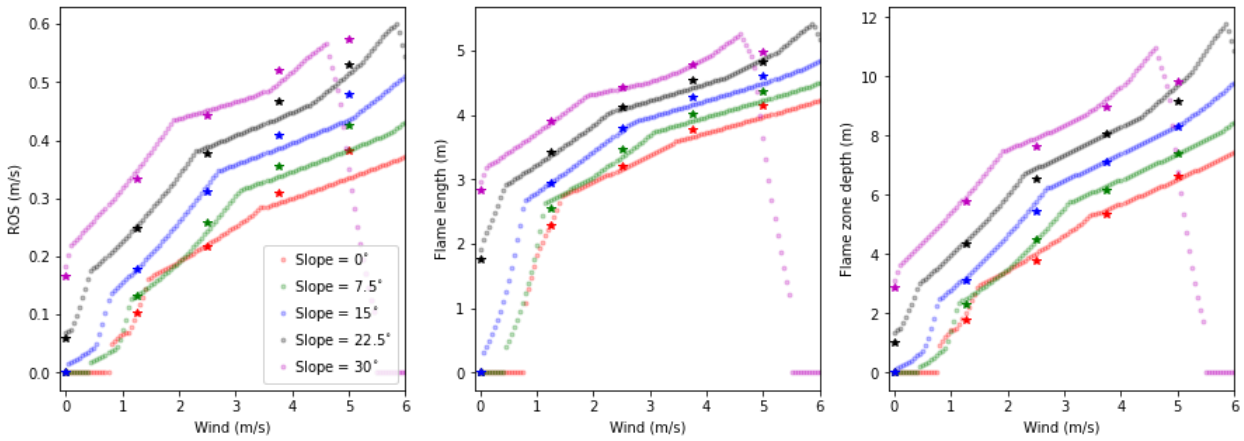

Figure 3: An example of the fit of the DL model (dots) to the 1D model outputs (stars with the same colour) for a range of wind speeds and slopes (0 to 30 degrees), with all other five input variables fixed at their corresponding median values. Extrapolated values are beyond the 5 m/s wind speeds.

4. Discussion

The current demonstration of DL modelling is limited to *homogeneous* fuel, weather, and terrain conditions, which is an oversimplification of the *heterogeneous* conditions where conditions include spatial varying fuels, time varying wind speed, etc. The physical 1D model can deal with this heterogeneity and future work involves investigating the neural network structure that would effectively model such conditions. FFNNs may no longer work as well, but it is likely that more advanced convolutional networks would suffice as demonstrated in Hodges and Lattimer (2019) and Radke et al. (2019).

5. Conclusions

We demonstrated an approach to wildfire modeling that relies on a physical fire spread model in reduced dimension (1D) to resolve the physical processes and DL modeling to produce a computationally efficient representation of the fire behavior. The DL model offers potential implementation in wildfire models at higher

dimensions (2D) such as FARSITE (Finney 1998) or 3D where computationally cheap methods are strongly desired. In this capacity, our DL model would perform with low cost compared to the original and be able to address behaviors resulting from fine-scale processes. The DL model accurately associates these steady state variables to the range of input conditions for fuels, weather, and topography, with mean (std) absolute errors of 0.078 (0.17) m/s, 0.21 (0.29) m and 0.98 (1.53) m for ROS, FL and FZD respectively. DL modeling is thus demonstrated as a promising means to represent fine-scale fire behaviors for broad-scale application while maintaining adequate fidelity to the detailed physical processes and inputs. Future work includes modeling heterogeneous input variables and explicitly embedding the DL model in subsequent more sophisticated fire models.

6. Acknowledgements

We thank Zack Ontiveros, Jared Sisk and Carla Bromberg for their helpful discussions, and Lily Hu and Rif A. Saurous for reviewing and approving this manuscript. Isaac Grenfell and Bill Chatham contributed to the 1D model and training data.

7. References

Clark, T.L., Griffiths, M., Reeder, M.J. and Latham, D., 2003. Numerical simulations of grassland fires in the Northern Territory, Australia: A new subgrid-scale fire parameterization. *Journal of Geophysical Research: Atmospheres*, *108*(D18).

Cohen, J.D. and Finney, M.A., 2022a. Fuel particle heat transfer part 1: convective cooling of irradiated fuel particles. *Combustion Science and Technology*, pp.1-27.

Cohen, J.D. and Finney, M.A., 2022b. Fuel Particle Heat Transfer Part 2: Radiation and Convection during Spreading Laboratory Fires. *Combustion Science and Technology*, pp.1-27.

Dillon, Joshua V., Ian Langmore, Dustin Tran, Eugene Brevdo, Srinivas Vasudevan, Dave Moore, Brian Patton, Alex Alemi, Matt Hoffman, Rif A. Saurous. TensorFlow Distributions. arXiv preprint arXiv:1711.10604, 2017.

Finney, M.A., 1998. *FARSITE, Fire Area Simulator--model development and evaluation* (No. 4). US Department of Agriculture, Forest Service, Rocky Mountain Research Station.

Finney, M.A., McAllister, S.S., Forthofer, J.M. and Grumstrup, T.P., 2021. *Wildland Fire Behaviour: Dynamics, Principles and Processes*. CSIRO PUBLISHING.

Fons, W.L., 1946. Analysis of fire spread in light forest fuels. *Journal of Agricultural Research. 72 (3): 92-121*, *72*(3), pp.92-121.

Forthofer, J.M., M.A. Finney, A. Trouve, and M. Ahmad. LIHTFire: a high-resolution physics-based fire spread model in one dimension *This conference*.

Goodfellow I, Bengio Y, Courville A (2016) 'Deep Learning.' (MIT Press)

Hodges JL, Lattimer BY (2019) Wildland Fire Spread Modeling Using Convolutional Neural Networks. Fire Technology55, 2115–2142. doi:10.1007/s10694-019-00846-4.

Piyush Jain, Sean C.P. Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Steve Taylor, and Mike D. Flannigan. (2020) A review of machine learning applications in wildfire science and management. Environmental Reviews. 28(4): 478-505. https://doi.org/10.1139/er-2020-0019

Radke D, Hessler A, Ellsworth D (2019) FireCast: Leveraging Deep Learning to Predict Wildfire Spread. https://glovis.usgs.gov/app?fullscreen=0

Scott, J.H. and Burgan, R.E., 2005. Standard fire behavior fuel models: a comprehensive set for use with Rothermel's surface fire spread model. *Gen. Tech. Rep. RMRS-GTR-153. Fort Collins, CO: US Department of Agriculture, Forest Service, Rocky Mountain Research Station. 72 p.*, *153*.

Sullivan, A.L., 2009. Wildland surface fire spread modelling, 1990–2007. 1: Physical and quasi-physical models. *International Journal of Wildland Fire*, *18*(4), pp.349-368.