
Prediction of Boreal Peatland Fires in Canada using Spatio-Temporal Methods

Shreya Bali ^{*1} Sydney Zheng ^{*1} Akshina Gupta ^{*1} Yue Wu ^{*1} Blair Chen ^{*1} Anirban Chowdhury ^{*1}
Justin Khim ^{†2}

Abstract

Peat fires are the largest fires on earth in terms of fuel consumption and are responsible for a significant portion of global carbon emissions. Predicting fires in the peatlands can help decision makers and researchers monitor and prevent peat fires. Despite this, research on predicting peatland fires remains largely understudied as compared to prediction of other forms of fires. However, peatland fires are unique among fires and therefore require datasets and architectures attuned to their particular characteristics. In this paper, we present a new dataset, PeatSet, designed specifically for the problem of peatland fire prediction. In addition, we propose several models to tackle the problem of fire prediction for the peatlands. We develop novel neural architectures for peatland fire prediction, PeatNet and PT-Net, with a graph-based and a transformer-based architecture, respectively. Our results indicate that these new deep-learning architectures outperform a regression baseline from existing peatland research. Among all the tested models, PT-Net achieves the highest F1 score of 0.1006 and an overall accuracy of 99.84%.

1. Introduction

Peatlands are a type of wetland that include marshes, bogs, fens, and swamps. They sequester more than twice as much carbon as stored in the world’s forests despite covering only 3% of the Earth’s land area (International Union for Conservation of Nature, 2017; Turetsky et al., 2015). Climate change has exacerbated the magnitude and frequency of fires and the length of the fire season (Flannigan et al., 2009). Additionally, peat fires release a large amount of the carbon sequestered in peatlands, emitting massive amounts

of carbon dioxide. The combination of anthropogenic climate change and peat fires form a positive feedback loop of peatland carbon emissions.

Several features of peatland fires differentiate them from commonly-explored forest fire approaches. First, peat fires produce less heat than typical fires. Second, peat fires can occur underground, in wet or cold areas, and even under snow as exemplified by western Canadian peat fires (Thompson, 2020). Third, peat fires can last for months when they are smouldering underground, challenging the use of burn duration as a heuristic for fire severity. Fourth, factors such as soil carbon and soil moisture are far more important for peatland fires than forest fires since the soil itself is a fuel source. Consequently, normal fire prediction techniques are less effective on peatland fires.

Prior work on regular forest fire and peatland fire prediction do present baselines to test the suggested models against. Most notably, the current state-of-the-art machine learning algorithms for forest fire prediction use convolutional neural networks (CNN) (Hodges et al., 2019; Radke et al., 2019).

Within the peatland research domain, there has only been a handful of studies in fire prediction, and far fewer that leverage deep-learning methods. Honma et al. (2016); Listyorini and Rahim (2018) use a system of detectors near a specific peatland to predict fire spread. Bourgeau-Chavez et al. (2020a) perform a review of four fires in peatlands to determine how the type of a peatland affects its likelihood of burning (Bourgeau-Chavez et al., 2020b). Maulana et al. (2019) use logistic regression to predict active fire areas. However, these models fail to scale well and are unable capture the complex relation between different causes as well as needed. Due to space constraints, we defer a more detailed discussion of previous work in Appendix I.

We propose bridging between the technical research on deep learning architectures and the scientific research on peatland fires to create a powerful prediction model for peatland fires. In line with this, our first general contribution is a new dataset, PeatSet, designed specifically for the problem of peatland fire prediction using previously existing datasets in the region of Canada. Our second main contribution, are two novel neural architectures for peatland fire prediction, PeatNet and PT-Net that we propose, implement, and

^{*}Equal contribution [†] The contributions in this paper were made prior to joining Amazon. ¹Carnegie Mellon University ²Amazon. Correspondence to: Shreya Bali <balishreya1@gmail.com>.

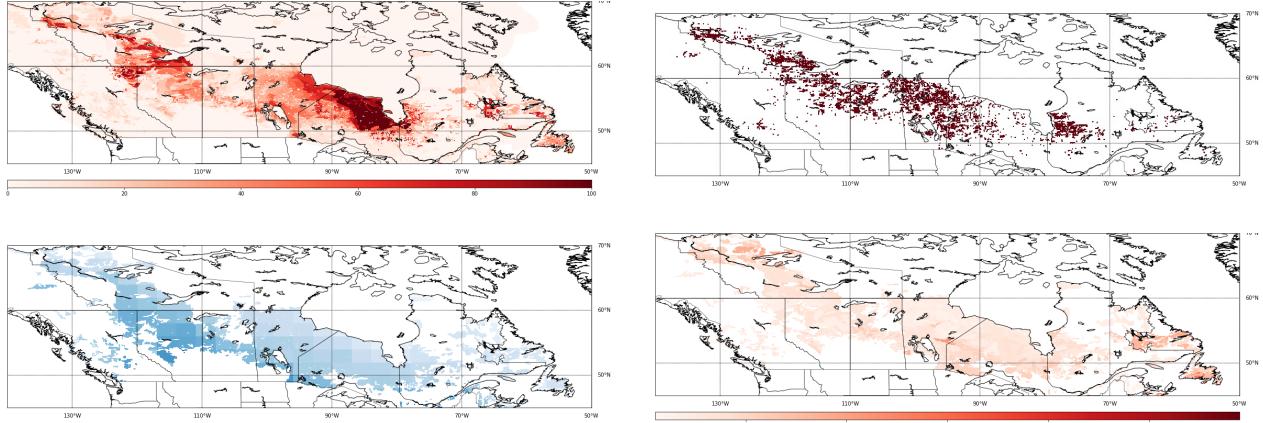


Figure 1. (upper-left) Tarnocai Peatland heatmap indicating percentage(%) of peatland cover. (upper-right) Burned area in the peatlands by masking the CWFIS burned area with Tarnocai. CWFIS data was max aggregated from 2012 to 2018. (lower-left) CO₂ emissions in the peatlands by masking the Global CarbonTracker data with Tarnocai. Data is in units of molar density, and is mean aggregated from 2012 to 2018. (lower-right) Total organic carbon content (kg/m²) in the peatlands by the Tarnocai dataset.

test. PeatNet is a graph-based architecture, and PT-Net is a transformer-based architecture.

This paper is organized as follows: in section 2, we highlight our dataset and model contributions; in section 3, we present our results in Results and in section 4, we present an analysis of the results .

2. Methodology

In this section, we elaborate on our two main contributions: the dataset and the models.

2.1. Dataset: PeatSet

Our first contribution is the curation of the first comprehensive peat fire dataset, PeatSet, consisting of previously existing remote sensing and manually labelled datasets. The spatial region of our dataset covers Canada because of the large area of peatlands and its relative abundance of publicly available data. We use PeatSet for the tasks of predicting CWFIS burned area categories. Figure 1 presents visualizations of the key features used over the Canadian region. Table 2 in Appendix II presents the features used for the prediction of fires.

Peatland Features: To delineate peatland from other land, we use the Tarnocai Peatland Map, which is the standard dataset used for determining where peatlands are in Canada. The map is divided into polygons, where each has an associated percentage of peatland cover, PEATLAND_P.

Fire Features: We use the burned area product from the Canadian Wildland Fire Information System (CWFIS) (Service, b), the most comprehensive dataset for fires in Canada.

The data is in part manually reported from governmental agencies, and should have greater inclusion of the unique kind of fires in peatlands as compared to remote sensing sources.

CO₂ Emissions Features: We use the Global Monitoring Laboratory Carbon Tracker CT2019, which has a 3-by-2 longitude/latitude resolution across the globe for CO₂ emissions. CT2019 also includes the flux of CO₂ across the globe. Flux is the gradient of concentration, and determines the source of the CO₂ emissions and the cause, e.g. fire, fossil fuels. CO₂ emissions is an indirect measurement of the presence of underground fires, as underground fires output significant CO₂.

Soil Features: We include features pertaining to the amount of carbon stored in the land, given by the Tarnocai dataset. Store carbon is a basic indicator of how much CO₂ is emitted if a fire burns over a given area.

Hot Spot Features: We incorporate hot spot data from the VIIRS dataset, as it indicates where fires are burning. With a finer resolution, the VIIRS satellite sensor imagery captures smaller fires as compared to MODIS, the standard satellite-based data used by many fire datasets such as the Global Fire Emissions Database (GFED). In general, hot spots have a confidence measurement of being associated with a fire, based upon its temperature. Since peatland fires burn at a lower temperature than traditional fires, low-confidence hot spots that persist over a long duration may still indicate the presence of peatland fires.

Additional Soil and Weather Features: From ERA5, we use soil moisture and soil temperature as indicators of a fire and wind velocity to account for the spread of CO₂ from a

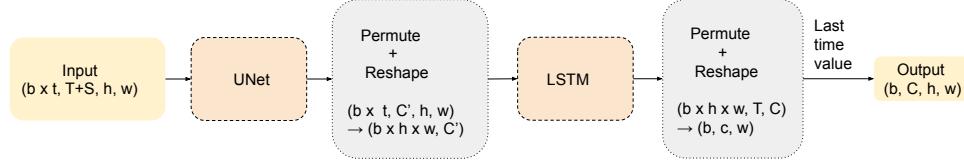


Figure 2. UNet-LSTM Architecture: Each batch is first fed into a U-Net component and then reshaped before being passed into the LSTM layer. b : batch size, h : height, C' : number of output channels from the U-Net component, C : number of output channels where there is one for each class.

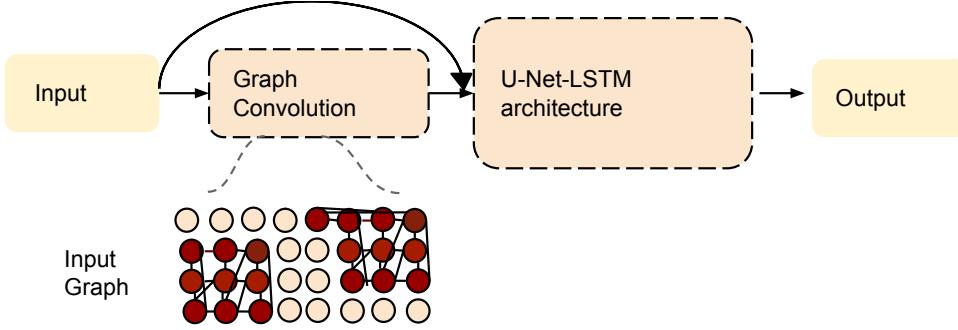


Figure 3. PeatNet Architecture: Dark points represent peatland, and light points represent non-peatland. Each node also has a self-loop in addition to the edges shown. Refer to the UNet-LSTM diagram in Figure 2 for a detailed description of its architecture.

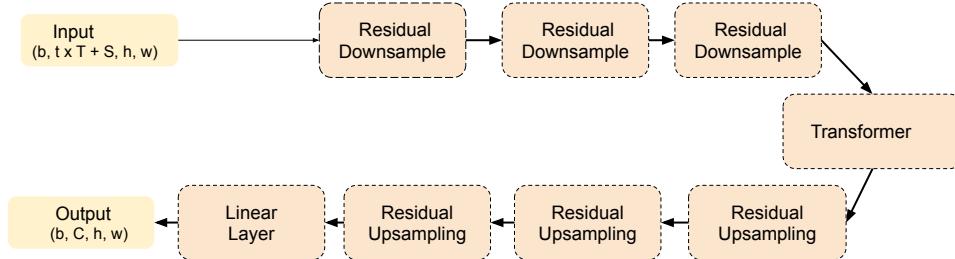


Figure 4. PT-Net Architecture: The input batch is encoded by the down-sampling layers, accumulated by the attention modules, and then up-sampled into the output.

source.

Refer to Appendix II for further details on the dataset features and processing.

2.2. Models

We test several models including U-Net, UNet-LSTM, PeatNet, and PT-Net. In order to compare our results to previous work, we also implement and test a logistic regression model as a baseline. This section first describes an auxiliary model we use, UNet-LSTM, and the two novel model architectures used for the prediction task, PT-Net and Peatnet.

2.3. UNet-LSTM

We implement UNet-LSTM, based on the U-Net model (Ronneberger et al., 2015), which is able to learn both the spatial information and temporal information. The diagram for the UNet-LSTM is shown in Figure 2. The inputs are first passed through the U-Net component and then to the LSTM layer. The U-Net component considers the temporal information as a part of the batch dimension, and thus only learns the spatial features. The output of this component is then reshaped such that the spatial features are a part of the batch dimension and the temporal features no longer are; this is then fed into the LSTM layer, which learns only temporal relations.

	Recall	Precision	F1	Accuracy
LR	0.8186	0.0016	0.0032	0.4298
U	0.9906	0.0212	0.0419	0.9607
UL	0.9944	0.0294	0.0571	0.9650
PN	0.9668	0.0274	0.0532	0.9632
PT	0.9232	0.0532	0.1006	0.9984

Table 1. Prediction results. Accuracy is high across all the models since there are more non-fire points as compared to fire points. It is therefore more meaningful to examine the recall, precision, and F1. Here, PT-Net dominates the other models. LR - Logistic Regression, U - UNet, UL - UNet-LSTM, PN - PeatNet, PT - PT-Net.

2.3.1. PEATNET

We first propose PeatNet, a novel graph network model, which is a graph-based neural network model designed to account for the underground spread of peat fires. Refer to Figure 3 for the outline of the architecture. The landmass is represented as being composed of a grid of points, where each point is a node in the graph. Nodes that model a point on peatland are connected to other peatland nodes within a distance of k , a hyperparameter; though a fire may spread far away, it is unlikely to spread past a certain distance. Each node also has a self-edge to model that a site with a fire will likely continue to have fire. A graph convolution is then applied to the graph, such that the nodes representing peatland gain information about other peatland nodes. We then pass the graphical convolution result and the original input features to a UNet-LSTM component. Finally, the output of the UNet-LSTM is passed through a fully connected layer to yield the final output.

2.3.2. PT-NET

The second model that we propose is PT-Net. Recent state-of-the-art studies show that attention-based neural networks are able to capture sequence-based data better than other neural networks, such as Long Short-Term Memory (LSTM) models. Our model is based on a residual encoder-decoder (He et al., 2015) with three down-sampling blocks in the encoder and three up-sampling blocks in the decoder, as shown in Figure 4. We account for temporal relations with a transformer module (Vaswani et al., 2017), which has three multi-head self-attention layers that focus on multiple past time-steps to predict the future time-step.

3. Results

In this section, we detail the training and testing and present the results in Table 1.

The models are evaluated on recall, precision, F1, and accuracy. Each model is provided 5 days of covariate input

data and predicts CWFIS burn classes for the subsequent day. The data is split into 70% train, 15% test and 15% validation days.

Note that for this task, recall is defined as the fraction of the fires correctly predicted over the total number of fires. Precision is defined as the fraction of the correct fires predicted over the total number of fires predicted.

We use binary cross entropy as the loss function for training the models. Note that the dataset is heavily skewed since there are far more non-fire points than fire points. We weigh the fire class a thousand times more heavily than the non-fire class because there are about one thousand non-fire data-points in training to a fire data-point.

4. Discussion

We first assemble a collection of relevant datasets to enable future studies of peatland fires. We hope that the data collection we provide will facilitate further research into peatland fire prediction. Additionally, we develop several novel architectures and adapt recent machine learning models to the problem of peat fire prediction. Our best model, PT-Net, shows great improvement in performance over previous models for fire prediction.

Our experiments show that models that consider spatiotemporal aspects of the data outperform those that do not. The regression model, which does not have information from nearby blocks, performs drastically worse than the models with access to spatial information. The UNet-LSTM and PT-Net, which both use spatial and temporal features, outperform the U-Net, which does not capture temporal information. PT-Net outperforms U-Net, UNet-LSTM, and PeatNet. A likely explanation for the higher performance of UNet-LSTM as compared to PeatNet is that locality generally dominates in fire modeling, and therefore, the long-distance relationships PeatNet captures through its graph-layer are less relevant.

This work can easily be expanded to other peatlands across the world with appropriate dataset expansion. The flexibility of our neural network approach allows additional features to be easily integrated. Additionally, accurate fire spread and severity prediction can allow decision makers to invest their attention to peatlands at high risk for fires and take appropriate preventative actions. By applying the tools developed in this work to the fire safety industry, we can significantly mitigate the carbon emissions that contribute to climate change today and reduce the damage caused by peat fires, as well as preserve existing peatland ecosystems.

References

- A. R. Jacobson, K. N. Schuldt, J. B. M. T. O. P. T. A. A. J. M. L. O. G. J. C. T. A. S. A. K. A. S. A. F. A. B. B. P. B. A. B. S. C. B. A. B. D. B. G. B. J. B. A. G. C. H. C. L. C. . M. Z. (2020). Carbontracker ct2019. model published by noaa earth system research laboratory, global monitoring division, 2020.
- Bourgeau-Chavez, L. L., S. Endres, R. Powell, M. J. Battaglia, B. Benscoter, M. Turetsky, E. S. Kasischke, and E. Banda (2017). Mapping boreal peatland ecosystem types from multitemporal radar and optical satellite imagery. *Canadian Journal of Forest Research* 47(4), 545–559.
- Bourgeau-Chavez, L. L., S. L. Grelak, M. Billmire, L. K. Jenkins, E. S. Kasischke, and M. R. Turetsky (2020a). Assessing boreal peat fire severity and vulnerability of peatlands to early season wildland fire.
- Bourgeau-Chavez, L. L., S. L. Grelak, M. Billmire, L. K. Jenkins, E. S. Kasischke, and M. R. Turetsky (2020b). Assessing boreal peat fire severity and vulnerability of peatlands to early season wildland fire.
- Castelli, M., L. Vanneschi, and A. Popović (2015). Predicting burned areas of forest fires: an artificial intelligence approach. *Fire ecology* 11(1), 106–118.
- Coffield, S. R., C. A. Graff, Y. Chen, P. Smyth, E. Foufoula-Georgiou, and J. T. Randerson (2019). Machine learning to predict final fire size at the time of ignition. *International journal of wildland fire*.
- Connolly, J. and N. M. Holden (2017). Detecting peatland drains with object based image analysis and geoeye-1 imagery. *Carbon balance and management* 12(1), 7.
- Cortez, P. and A. d. J. R. Morais (2007). A data mining approach to predict forest fires using meteorological data.
- DeLancey, E. R., J. Kariyeva, J. T. Bried, and J. N. Hird (2019). Large-scale probabilistic identification of boreal peatlands using google earth engine, open-access satellite data, and machine learning. *Plos one* 14(6), e0218165.
- Flannigan, M. D., M. A. Krawchuk, W. J. de Groot, B. M. Wotton, and L. M. Gowman (2009). Implications of changing climate for global wildland fire. *International journal of wildland fire* 18(5), 483–507.
- Ganapathi Subramanian, S. and M. Crowley (2018). Using spatial reinforcement learning to build forest wildfire dynamics models from satellite images. *Frontiers in ICT* 5, 6.
- He, K., X. Zhang, S. Ren, and J. Sun (2015). Deep residual learning for image recognition. *corr abs/1512.03385* (2015).
- Hersbach, H., B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al. (2020). The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* 146(730), 1999–2049.
- Hodges, J. L., B. Y. Lattimer, and K. D. Luxbacher (2019). Compartment fire predictions using transpose convolutional neural networks. *Fire Safety Journal* 108, 102854.
- Honma, T., K. Kaku, A. Usup, and A. Hidayat (2016). Detection and prediction systems of peat-forest fires in central kalimantan. In *Tropical Peatland Ecosystems*, pp. 397–406. Springer.
- Hugelius, G., J. Loisel, S. Chadburn, R. B. Jackson, M. Jones, G. MacDonald, M. Marushchak, D. Olefeldt, M. Packalen, M. B. Siewert, et al. (2020). Large stocks of peatland carbon and nitrogen are vulnerable to permafrost thaw. *Proceedings of the National Academy of Sciences* 117(34), 20438–20446.
- International Union for Conservation of Nature (2017, November). Peatlands and climate change.
- Jin, G., C. Zhu, X. Chen, H. Sha, X. Hu, and J. Huang (2020). Ufsp-net: a neural network with spatio-temporal information fusion for urban fire situation prediction. In *IOP Conference Series: Materials Science and Engineering*, Volume 853, pp. 012050. IOP Publishing.
- Kalacska, M., J. P. Arroyo-Mora, R. J. Soffer, N. T. Roulet, T. R. Moore, E. Humphreys, G. Leblanc, O. Lucanus, and D. Inamdar (2018). Estimating peatland water table depth and net ecosystem exchange: A comparison between satellite and airborne imagery. *Remote Sensing* 10(5), 687.
- Liang, H., M. Zhang, and H. Wang (2019). A neural network model for wildfire scale prediction using meteorological factors. *IEEE Access* 7, 176746–176755.
- Listyorini, T. and R. Rahim (2018). A prototype fire detection implemented using the internet of things and fuzzy logic. *World Trans. Eng. Technol. Educ* 16(1), 42–46.
- Lozhkin, V., D. Tarkhov, V. Timofeev, O. Lozhkina, and A. Vasilyev (2016). Differential neural network approach in information process for prediction of roadside air pollution by peat fire. In *IOP conference series: materials science and engineering*, Volume 158, pp. 012063. IOP Publishing.

- Mahdianpari, M., B. Salehi, M. Rezaee, F. Mohammadi-mansh, and Y. Zhang (2018). Very deep convolutional neural networks for complex land cover mapping using multispectral remote sensing imagery. *Remote Sensing* 10(7), 1119.
- Markuzon, N. and S. Kolitz (2009). Data driven approach to estimating fire danger from satellite images and weather information. In *2009 ieee applied imagery pattern recognition workshop (aipr 2009)*, pp. 1–7. IEEE.
- Maulana, S. I., L. Syaufina, L. B. Prasetyo, and M. N. Aidi (2019). Spatial logistic regression models for logistic regression models for predicting peatland fire in bengkalis regency, indonesia. *Journal of Sustainability Science and Management* 14(3), 55–66.
- Minasny, B., B. I. Setiawan, S. K. Sapomo, A. B. McBratney, et al. (2018). Open digital mapping as a cost-effective method for mapping peat thickness and assessing the carbon stock of tropical peatlands. *Geoderma* 313, 25–40.
- Mitsopoulos, I. and G. Mallinis (2017). A data-driven approach to assess large fire size generation in greece. *Natural Hazards* 88(3), 1591–1607.
- Radke, D., A. Hessler, and D. Ellsworth (2019). Forecast: Leveraging deep learning to predict wildfire spread. In *IJCAI*, pp. 4575–4581.
- Ronneberger, O., P. Fischer, and T. Brox (2015). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer.
- Service, C. F. Canadian national fire database – agency fire data. Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre, Edmonton, Alberta.
- Service, C. F. National burned area composite (nbac). Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre, Edmonton, Alberta.
- Shidik, G. F. and K. Mustofa (2014). Predicting size of forest fire using hybrid model. In *Information and Communication Technology-EurAsia Conference*, pp. 316–327. Springer.
- Sitanggang, I. S., R. Yaakob, N. Mustapha, and A. Ainuddin (2014). A decision tree based on spatial relationships for predicting hotspots in peatlands. *Telkomnika* 12(2), 511.
- Subramanian, S. G. and M. Crowley (2017). Learning forest wildfire dynamics from satellite images using reinforcement learning. In *Conference on reinforcement learning and decision making*.
- Tarnocai, C., I. M. Kettles, and B. Lacelle (2011a). Peatlands of Canada. Geological Survey of Canada.
- Tarnocai, C., I. M. Kettles, and B. Lacelle (2011b). Soil Organic Carbon Content of Canadian Peatlands. Geological Survey of Canada.
- Tarnocai, C., I. M. Kettles, and B. Lacelle (2011c). Soil Organic Carbon Mass of Canadian Peatlands. Geological Survey of Canada.
- Thompson, D. (2020, July). Peatland fires and carbon emissions.
- Turetsky, M. R., B. Benscoter, S. Page, G. Rein, G. R. Van Der Werf, and A. Watts (2015). Global vulnerability of peatlands to fire and carbon loss. *Nature Geoscience* 8(1), 11–14.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin (2017). Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008.
- Widyatmanti, W., D. Umarhadi, M. Ningam, Z. Sarah, K. Nugroho, Y. Sulaeman, et al. (2019). Mapping acid sulfate soil hydrogeomorphical unit on the peatland landscape using a hybrid remote sensing approach. In *Tropical Wetlands-Innovation in Mapping and Management: Proceedings of the International Workshop on Tropical Wetlands: Innovation in Mapping and Management, October 19-20, 2018, Banjarmasin, Indonesia*, pp. 30. CRC Press.
- Zheng, Z., W. Huang, S. Li, and Y. Zeng (2017). Forest fire spread simulating model using cellular automaton with extreme learning machine. *Ecological Modelling* 348, 33–43.

Appendix I: Previous Work

Forest and Urban Fire Prediction

Fire behavior prediction methods generally focus on predicting growth and spread or predicting final severity. In order to predict severity, many models use metrics such as final burned area or duration of burn. Using the former assumes severity increases when more land is burned, and using the latter assume severity increases the longer the fire lasts.

Common fire behavior prediction methods are regression (Castelli et al., 2015; Cortez and Morais, 2007; Mitsopoulos and Mallinis, 2017), random forests (Markuzon and Kolitz, 2009; Mitsopoulos and Mallinis, 2017), support vector machines (Castelli et al., 2015; Cortez and Morais, 2007), or Bayesian networks (Markuzon and Kolitz, 2009).

Researchers have traditionally formulated the fire prediction problem as a classification problem. The highest accuracy among these models is 97.5% and is achieved by Shidik and Mustofa (2014). Similar research achieves far lower accuracy (Coffield et al., 2019; Mitsopoulos and Mallinis, 2017). However, these lower-accuracy works create classes based upon the ground-truth burned area size instead of clusters on the covariates as done by Shidik and Mustofa (2014).

While few works have explicitly accounted for temporal information, Liang et al. (2019) compare a recurrent neural network (RNN) and an LSTM to predict a numerical custom fire severity metric. Their results indicate that the LSTM outperforms the RNN, motivating our use of an LSTM. In addition, they discuss that meteorological covariates, which they use, are associated with fire severity. These features may therefore be worth considering in the peat fire detection and prediction problem as well.

However, a two-dimensional map of predicted fire perimeters is easier for researchers, policymakers, and fire agencies to analyze and use. Recently, deep learning methods such as reinforcement learning, CNNs, and graph neural network (GNN) models have gained more attention in mapping fire spread.

Reinforcement learning models produce predicted fire perimeters by viewing the fire as an agent and modelling actions the agent is likely to take (Zheng et al., 2017; Subramanian and Crowley, 2017; Ganapathi Subramanian and Crowley, 2018). Typically, the fire is modelled so it can only move to nearby areas in a timestep. However, this formulation of the problem does not apply in the context of peatland fires since fires are no longer limited to spreading to the areas around them. Peat fires can go underground and resurface elsewhere, which is parallel to a delayed jump action for a fire agent. As such, traditional reinforcement learning algorithms would have to be altered in order to

function for peatlands.

The current state-of-the-art machine learning algorithms for forest fire prediction use convolutional neural networks (CNN) (Hodges et al., 2019; Radke et al., 2019). Hodges et al. (2019) to predict future burn perimeters based on six-hourly burn maps generated by the FARSITE physics-based simulator. Radke et al. (2019) attempts to use a similar CNN architecture named FireCast based on daily observed fire perimeters from GEOMAC instead of simulation burns. FireCast is able to outperform FARSITE, which establishes some of the limitations of using simulated burns for training as compared to observed data. FireCast's performance emphasizes recall over precision with very high recall percentages but very low precision.

Few works take into account abnormally shaped spatial information for fire prediction. However, Jin et al. (2020) uses graph-convolutional layers in a custom architecture, USFP-Net. The USFP-Net uses a graph convolutional neural network, CNN layers and RNNs to model the fire prediction problem. USFP-Net outperforms many other common architectures for urban fire prediction. They represent the area as a graph with edges with weights inversely proportional to the distance between them, resulting in a fully connected graph. Training such a network would require a long time, and it is likely computationally infeasible to use this model/replicate such results over larger areas. In addition, the model only accounts for the usual spatial and temporal characteristics of an area without taking into account fire spread patterns.

Peatland Studies

We have surveyed literature within the general peatland domain that is potentially relevant to the problem of peatland fire prediction. Many problems arise in characterizing the peatland biome. First, it can be difficult to determine which land is peatland. DeLancey et al. (2019) predict where peatlands exist by using machine learning and boosted decision trees. Mahdianpari et al. (2018) classify wetlands as being peat-based or not by using a convolutional neural network model. Second, the peatlands have various characteristics that might be useful as covariates in fire prediction. Prior work on these elements focus on identifying the type of peatlands (Bourgeau-Chavez et al., 2017), the amount of sequestered carbon stored in the peat (Minasny et al., 2018), peatlands affected by permafrost (Hugelius et al., 2020), the acidity of the peatlands (Widyatmanti et al., 2019), identifying human draining around peatlands (Connolly and Holden, 2017), or the water table depth of the peatlands (Kalacska et al., 2018) using various remote-sensing datasets, statistical methods, or basic machine learning models.

Peatland Fire Prediction

Within the peatlands domain, there has only been a handful of studies in fire prediction, and far fewer that leverage deep-learning methods.

[Honma et al. \(2016\)](#); [Listyorini and Rahim \(2018\)](#) use a system of detectors near a specific peatland to predict fire spread. [Bourgeau-Chavez et al. \(2020a\)](#) perform a review of four fires in peatlands to determine how the type of a peatland affects its likelihood of burning ([Bourgeau-Chavez et al., 2020b](#)). However, these works have limited scalability since they are either restricted to areas in which they can establish a detection system or are restricted to post-event analysis.

[Sitanggang et al. \(2014\)](#) apply decision trees to determine active fire area, with their best model reaching an accuracy of 71.66%. However, this model does not perform any forward predictions; it predicts MODIS hot spots based upon the conditions of that very day. This work therefore shows that hot spots can be correlated with various climate and landscape information such as soil moisture, vegetation type, and precipitation.

[Maulana et al. \(2019\)](#) use logistic regression to predict active fire areas up to three months in advance, achieving an accuracy of 85.16%. The authors average climate and landscape conditions for four months to predict MODIS fire hot spots aggregated for one month.

[Lozhkin et al. \(2016\)](#) use a differential neural network model to predict carbon monoxide dispersion from peat fires near highways. This research suggests the ability of neural networks to capture peat emissions patterns.

Appendix II: Dataset

Feature Processing: The time range of our dataset is from January 20, 2012 to December 31, 2018, which is the intersection of the available time ranges for the different features. We use the following South/North/West/East coordinates to bound the map for each feature: -141.0000°/-50.0000°/41.7500°/90.0000°. The maps are explicitly projected to WGS84 (EPSG: 4326), the standard longitude/latitude coordinate system. As each dataset has a different spatial resolution, they are all scaled to have a resolution of 0.1° x 0.1° longitude/latitude per pixel, such that each feature at a timestamp has dimensions of 483 by 910 pixels. We take timestamps per day. For features with a sub-daily resolution, we take the average over sub-daily data points. For features that do not change over time, such as TOCC from the Tarnocai dataset, we simply reuse the same data for each timestamp.

Additional Predicted Feature Processing: For prediction, we mask the predicted features, fire occurrence or CO₂ emis-

sions, with the Tarnocai Peatland shapefiles to get only the values over peatlands; we do not do this for the covariates. A polygon in the Tarnocai Peatland Map is considered to be peatland if it has at least 10% peatland cover as specified by PEATLAND_P. A polygon in CWFIS is considered to be burned if it is estimated 100% burned as given by the BURNCLAS feature.

We list features we considered interesting from the following datasets, whether they were eventually used or not.

1. [Tarnocai Peatland Map \(Tarnocai et al., 2011a\)](#)([Tarnocai et al., 2011b](#))[\(Tarnocai et al., 2011c\)](#)

A set of shapefiles capturing where the peatlands are, what type (bog, fen, swamp marsh) they are, and how much carbon they store. The data was gathered through survey and published by National Resources Canada.

(a) Spatial Range: Canada

(b) Variables:

- i. Peatland (%) (PEATLAND_P): Percentage of shapefile polygon covered in peatland.
- ii. BOG_PCT: Percentage of shapefile polygon covered in bog.
- iii. FEN_PCT: Percentage of shapefile polygon covered in fen.
- iv. SWAMP_PCT: Percentage of shapefile polygon covered in swap.
- v. MARSH_PCT: Percentage of shapefile polygon covered in marsh.
- vi. TOCC: The average amount of carbon stored per surface area (kg/m²) for the shapefile polygon (uses peat depth).

2. [Canadian Wildland Fire Information System \(CWFIS\) \(Service, b\)](#)([Service, a](#))

Shapefiles indicating where fires occurred and their burned areas. Data on fires was collected through survey by Canadian fire management agencies. Data on burned area was calculated through a combination of survey data, aerial photography, and satellite data, such as from through MODIS, VIIRS, Landsat, and Sentinel-2.

(a) Spatial Range: Canada

(b) Temporal Resolution: Daily

(c) Temporal Range: ; January 2000 - January 2019

(d) Variables:

- i. BURNCLAS: Proportion of land burned for shapefile polygon. 1: estimated 25% burned, 2: estimated 50% burned, 3: estimated 75% burned, 4: estimated 100% burned.

Features	Dataset	Spatial Resolution	Temporal Resolution
BURNCLASS	CWFIS	variable	daily
height_i (for i in $[0, 10)$)	CarbonTracker (Global)	$3^\circ \times 2^\circ$	3-hourly
fire_flux	CarbonTracker (Flux)	$1^\circ \times 1^\circ$	3-hourly
fuel_flux	CarbonTracker (Flux)	$1^\circ \times 1^\circ$	3-hourly
frp	VIIRS	375m x 375m	daily
confidence	VIIRS	375m x 375m	daily
bright_t14	VIIRS	375m x 375m	daily
TOCC	Tarnocai Peatland Map	variable	fixed
swvli (for i in $[1, 4]$)	ERA5	$0.1^\circ \times 0.1^\circ$	hourly
stli (for i in $[1, 4]$)	ERA5	$0.1^\circ \times 0.1^\circ$	hourly
lai_hv	ERA5	$0.1^\circ \times 0.1^\circ$	hourly
lai_lv	ERA5	$0.1^\circ \times 0.1^\circ$	hourly
tp	ERA5	$0.1^\circ \times 0.1^\circ$	hourly
t2m	ERA5	$0.1^\circ \times 0.1^\circ$	hourly
u10	ERA5	$0.1^\circ \times 0.1^\circ$	hourly
v10	ERA5	$0.1^\circ \times 0.1^\circ$	hourly

Table 2. Covariate features used for the prediction of fire. Datasets are presented along with the relevant covariate features, spatial resolution, and temporal resolution. Spatial resolution indicated with degrees is given by longitude/latitude.

ii. Fire: Shapefile polygons outlining fires.

3. VIIRS

Location of hot spots, which are thermal anomalies that often indicate fire. The data is gathered by the VIIRS sensor, onboard the Suomi NPP and NOAA-20 polar-orbiting satellites.

- (a) Spatial Range: Global
- (b) Spatial Resolution: 375m x 375m
- (c) Temporal Range: January 2012 - present
- (d) Temporal Resolution: Daily
- (e) Variables:
 - i. frp: Fire radiative power, megawatts.
 - ii. confidence: Confidence of individual hot spots. 0: low, 1: nominal, 2: high.
 - iii. bright_t14: Fire pixel channel I4 brightness temperature (Kelvin).

4. Global Monitoring Laboratory Carbon Tracker CT2019 Globe (A. R. Jacobson, 2020)

- (a) Spatial range: Global (NOTE: During winter time, high latitude regions are not reliable.)
- (b) Spatial resolution: 3° longitude x 2° latitude
- (c) Temporal Range : January 2000 - present
- (d) Temporal Resolution: 3-hourly
- (e) Variables:
 - i. height_i, for $i \in [0, 10)$: CO₂ molar density at 10 different height levels above the ground. Refer to Table 1 in Section 6.1 on

their documentation for the actual heights in meters.

5. Global Monitoring Laboratory Carbon Tracker CT2019 Flux (A. R. Jacobson, 2020)

- (a) Spatial range: Global (NOTE: During winter time, high latitude regions are not reliable.)
- (b) Spatial resolution: $1^\circ \times 1^\circ$
- (c) Temporal Range: January 2000 - present
- (d) Temporal Resolution: 3-hourly
- (e) Variables:
 - i. fire_flux: Flux of CO₂ attributed to fire.
 - ii. fuel_flux: Flux of CO₂ attributed to burning fossil fuel.

6. ERA5 (Hersbach et al., 2020)

The standard dataset on weather variables, such as pertaining to soil, precipitation, temperature, and wind. The data is an assimilation between observations and modelling of climate.

- (a) Spatial Resolution: 9km x 9km regridded for $0.1^\circ \times 0.1^\circ$
- (b) Temporal Resolution: hourly
- (c) Temporal Range: January 1981 - 3 months before present
- (d) Variables:
 - i. swvli: Soil water level 1. Meters-cubed water in meters-cubed soil at a depth of 0 - 7 cm from surface.

- ii. swv12: Soil water level 2. Meters-cubed water in meters-cubed soil at a depth of 7 - 28 cm from surface.
- iii. swv13: Soil water level 3. Meters-cubed water in meters-cubed soil at a depth of 28 - 100 cm from surface.
- iv. swv14: Soil water level 4. Meters-cubed water in meters-cubed soil at a depth of 100 - 289 cm from surface.
- v. st11: Soil temperature level 1. Temperature of soil in Kelvins at a depth of 0 - 7 cm from surface.
- vi. st12: Soil temperature level 2. Temperature of soil in Kelvins at a depth of 7 - 28 cm from surface.
- vii. st13: Soil temperature level 3. Temperature of soil in Kelvins at a depth of 28 - 100 cm from surface.
- viii. st14: Soil temperature level 4. Temperature of soil in Kelvins at a depth of 100 - 289 cm from surface.
- ix. lai_hv: Leaf area index, low vegetation. Surface area of low-lying leaves in meters-squared over area over area land in meters-squared. Characterizes the density of low vegetation such as crops, marshes, grasses, bogs.
- x. lai_lv: Leaf area index, high vegetation. Surface area of high-reaching leaves in meters-squared over area land in meters-squared. Characterizes the density of high vegetation such as forests and trees.
- xi. tp: Total precipitation. Total amount of water accumulated over an hour. Given as the depth in meters the water would have been if spread evenly over the spatial unit. CAUTION: This variable is an aggregation and not an average, so its value describes a very specific space and time.
- xii. t2m: 2-meter temperature. Temperature (K) two meters above the surface.
- xiii. u10: East-ward component of wind. Positive-x component of speed (m/s) of wind ten-meters above surface.
- xiv. v10: North-ward component of wind. Positive-y component of speed (m/s) of wind ten-meters above surface.