A Complementary Approach to Improve Wild Fire Prediction Systems

Sriram Ganapathi Subramanian
Electrical and Computer Engineering
University of Waterloo
Waterloo, ON, Canada N2L3G1
s2ganapa@uwaterloo.ca

Mark Crowley
Electrical and Computer Engineering
University of Waterloo
Waterloo, ON, Canada N2L3G1
mcrowley@uwaterloo.ca

Abstract

The physics based simulators used as the current state of the art in predicting wild fires have been found to be difficult to set up and operate, very slow, requiring large amounts of computational resources and high quality data sets, having an inherent bias that does not generalize and rendering low accuracies. In this regard there is a critical need for alternate tools that can augment or replace these physics based simulators. We demonstrate an approach that uses Reinforcement Learning for predicting the spread of wild fires over the next few days and/or months based on the the physical conditions of the landscape, the intensity and the location of the wild fire at present. The fire is modelled as the agent in this Markov Decision Process setting where it is designed to make a spread choice based on the physical conditions surrounding it. The agent learns to weight the relative importance of different environmental variables as it trains using copious amounts of data. This is a direct contrast to the physics based simulators where a large amount of expert rules are used to model the physical laws. We have run our suggested setting on real and synthetic fires and noted its specific advantages. We have also studied a particular fire simulation system and recognized its advantages and deficiencies. We propose that the approaches using recent advances in Artificial Intelligence techniques should be used in combination with the existing physics based simulators to entail more accurate and easy to use fire prediction models.

1 Introduction

Wild fires are becoming one of the most dangerous natural disasters in many countries. These are not seasonal in nature anymore and have become a perennial problem. The problem is escalating quickly. From the 1980s to 2015 the number of acres burned has grown from two million a year to eight million a year [1]. 2018 has particularly been a bad fire year, with wild fires raging in regions like British Columbia (Canada) [15] and California (United States) [3]. Wild fire prediction systems built on the physics based and/or rule based processes are the current state of the art in performing fire simulation and risk assessment. Many different types of fire simulators have been proposed and used in the wild fire modelling and prediction literature over the last few years. In-spite of this fact, wild fire researchers and practitioners have not been able to converge on a single comprehensive model of the fire behavior [16]. The authors in [11] evaluate some of the state of the art fire simulators using real wild fire datasets in Iran. The accuracy of most methods are reported to be around 50%. The major factors that must be improved for obtaining better accuracy estimates are reported to be availability of good quality input data, theoretical basis of the fire behaviour model and the fire growth algorithm which need to be improved. In [2], the authors survey the most popular fire behaviour models and conclude that there are very few or no fire behaviour models that can be easily used for risk assessment and fuel management in case of wild fires and the existing fire behaviour models all have distinct and complex data input and output formats, that are not easily available. The authors

32nd Conference on Neural Information Processing Systems (NIPS 2018), Montréal, Canada.

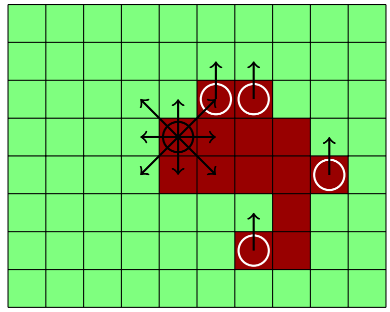

(a) Schematic of the state and actions

Figure 1: Forest Wildfire Satellite Data Domain: A schematic of the wildfire motion domain at a particular state and timestep. The red (dark) cells are on fire, the green (light) are not on fire, the dark circle indicates the current cell or agent being spread by the policy. The arrows around the dark circle indicates the action choices possible. The white circle indicates that other cells will be considered for spread. The arrows from the white circle indicates that there is a strong wind blowing towards North and the north action is the most likely action choice for these cells (effect of wind).

in [6] analyze contemporary fire modelling systems like FARSITE and NEXUS and demonstrate a significant under prediction bias in the simulations of fire modelling systems. The reasons for this bias in the designed expert rules of the systems are shown to come from incompatibility in model linkages, use of wrong fire rate spread models and uncalibrated custom fuel models. Thus, these models are build based on the most common fire and forest types in North America, which do not generalize well to all regions in North America and the rest of the world too. A wide range of sources for uncertainties in contemporary wild fire management systems in summarized in [26].

Our approach involves modelling the fire into an agent that determines the best direction to move based on the surrounding conditions. The agents explore/train using readily available real and synthetic fire data sets and do not depend on expert knowledge which makes it flexible with respect to the study area. The wildfire spreading problem is modelled into a Markov Decision Process (MDP), which is subsequently solved using existing and new Reinforcement Learning (RL) algorithms to get a fire spread policy. This approach is a "Top - Down" approach as the agent is given the complete fire burning scenario for the entire study area which it uses to learn the fire spread policy on individual pixels (small subsets) of the data. On the other hand the fire simulation systems in practice today have a "bottom up" approach where they have a predefined fire spread policy, that is applied to individual pixels to determine the overall burning areas in the study area given the ignition points. Upon initial experiments we have determined that the two approaches are complementary and have good potential to be combined with each other. An RL agent has been attempted in spatial process before [9] where the author states that MDPs and RL being commonly used for representing and solving sequential decision making problems under uncertainty is a good fit for tackling problems like forest fires. But the analysis stop at a simulation level itself. We have used real data sets from wild fire events to make a comprehensive analysis of the approach. Unsupervised approaches have also been tried in this domain where the results showed better performances than modern fire danger prediction systems in Australia [20]. This approach can take the dynamic nature of variables into account and work with unlabeled data, but it needs huge amounts of high resolution good quality data to be successful. Further, this only predicts the relative levels of risk of different regions in the study area, but does not provide a detailed fire spread policy to augment the existing systems like our proposed method.

2 Reinforcement learning Model and Experiments

Reinforcement learning [24] is an area of Machine Learning where the agent learns to act in an environment by exploring action choices in distinct states and receiving subsequent rewards. Rewards are not strong labels given to the agent, but a weak signal on how good or bad it did. The agent aims to learn an optimal policy that maximizes its rewards in the given scenario. Reinforcement Learning is being used in our research, for this domain due to several reasons. The first reason being that learning agents have never been used to understand the spread dynamics of complex natural processes using real world events. The second reason is that RL has recently shown superior performances in several domains and the wild fire domain seems like a natural extension to such domains with

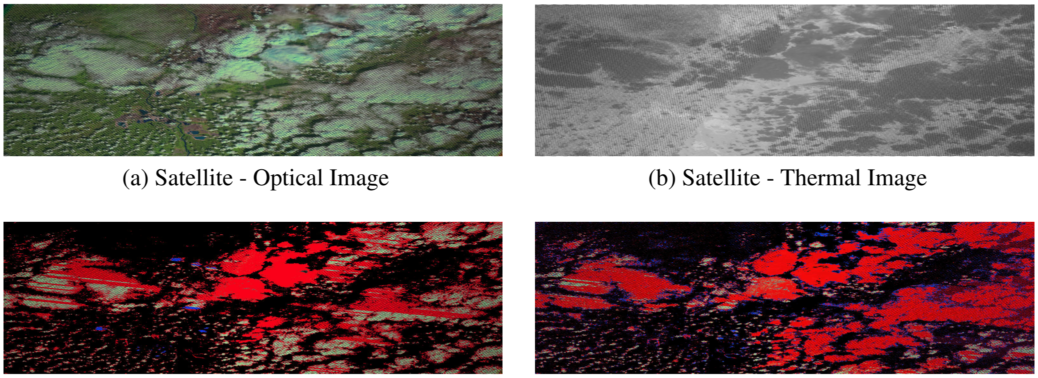

(a) Satellite - Optical Image (b) Satellite - Thermal Image

(c) A3C (d) MCTS-A3C

Figure 2: Results for an experiment that involve predicting forward in a wild fire setting. Performances of A3C and MCTS-A3C are shown along with the optical and thermal images from landsat.

increased complexity. For the forest fire spread problem, an RL approach can be a possible solution as the fire is influenced by the environment (direction of spread) and it also influences the environment (phenomenon of burning) which is a requirement in any RL set up [23].

Our model is described in some of our earlier publications [10] and [22]. We define a Markov Decision Process (MDP) $< S, A, P, R >$ where the set of states S describes any location on the landscape and where the 'agent' taking actions is a fire spreading across the landscape. A state $s \in S$ corresponds to the state of a cell in the landscape $(x, y, te, l, w, d, rh, r, i)$ where x, y are the grid coordinates of the cell, te is the temperature at the particular time and location and l is the land cover type of the cell derived from satellite images (values include: water, vegetation, built up, bare land, other), w is wind speed, d is wind direction, rh is relative humidity, i is the fireline intensity of fire at a cell as defined in [12], and r is the average amount of the rainfall at the spatial coordinates during the time of study. These variables are considered to be the most contributing factors to fire spread as they are the primary variables in the Canadian Forest fire weather index as specified in [5]. The action $a \in A$ indicates the direction the fire at a particular cell 'chooses' to move: North, North-West, North-East, South, South-West, South-East, East or West or to stay put. The reward function R maps a cell state to a continuous value in the range [-1,1]. We have two stages of reward assignment. The first stage is locally based and is contingent on the land cover of the cell on which the fire decides to spread. This function scores spread choices in proportion to flammability of a cell (Veg: 0.5, Water = -1, urban = -0.1, empty land = 0.05, Other = 0). The second stage is globally based on ground truth information. Figure 1 shows a schematic of the domain. The transition dynamics for the problem (for model based approaches) is derived from the physical dynamics of important state variables in this environment. The primary data sets for our problem came from medium resolution Landsat satellite images. The secondary data sources include Canada wildfire information portal and sources mentioned in [17].

Experiments using this model involved predicting the location of a fire a few days into the future given its current location, predicting the path of a fire and transferring the learned policy to a different fire in a similar region. We studied the performances of different types of RL algorithms and the results are summarized in [22]. We also tried a Gaussian Process supervised learning algorithm [19] to serve as a baseline. This was chosen as Gaussian Process have been known to perform well in problems with a spatial structure [4]. Tree based search techniques using roll out policy (Monte Carlo Tree search [13]) and actor critic based deep learning algorithms (A3C [14]) had particular advantages in different kinds on experiments. We also designed an algorithm, "MCTS-A3C" by a novel combination of these algorithms that had the unique advantages of both Tree Search and Actor Critic models [21]. Figure 2 shows some results that we obtained for a particular experiment that involves predicting forward the movement of fire given its ignition and landscape conditions. The results in figure 2 shows some chosen important algorithms, out of all the ones we tried. We have also displayed the corresponding satellite images from Landsat to give an idea of the performances. The red in the output image corresponds to true positives, the blue represents false positives, the white is false negatives and black is true negatives. For the same experiment in figure 2 the accuracy values obtained are summarized in table 1. The custom defined MCTS-A3C gives far better results on this experiment as compared

Algorithm	GP	VI	PI	QL	MCTS	A3C	MCTS-A3C
Avg Accuracy	50.8%	25.4%	38.2%	10.4%	60.2%	53.2%	90.8%

Table 1: Average Accuracy for each algorithm on an experiment involving predicting forward in a wild fire setting. The algorithms are Gaussian Processes (GP), Value Iteration (VI), Policy Iteration (PI), Q Learning (QL), Monte Carlo Tree search (MCTS), Asynchronous Advantage Actor Critic (A3C) and MCTS-A3C which is a custom designed [21].

to the other pre-existing state of the art RL algorithms. We have also performed experiments using simulated fire data sets and made comparisons of the RL approach to that of a popular wild fire simulation system (BP3 [18]) [21]. We have based our accuracy comparisons and reported results predominantly on real fire experiments.

3 Current and Future Work

With regards to our current research focus, we are making metric comparisons between the performances of our RL approach and additional state of the art wild fire simulator systems. The results with BP3 wild fire prediction system indicated that BP3 had advantages in experiments where the available data samples were error prone and sufficient data samples were not available. Further, we observed that BP3 did quite well in small scale experiments where the study area was in the order of a few kilometers. The RL approach did quite well in experiments with sufficient amounts of good quality data samples and in experiments were the study area is quite large (in the order of few thousand kilometers) [21]. Furthermore, for large study areas, BP3 which follows a cell by cell approach take a prohibitively large amount of time to execute completely. Thus, there is a clear indication here that a potential combination of the two distinct approaches, namely RL and BP3, would lead to a better prediction model overall.

The different wild fire simulation systems we are considering follows different simulation models as elaborated in [16]. In particular we are focusing on FARSITE [8] which is one of the most widely used physics based simulator for industrial applications. The algorithm that we are considering takes the form of having a global agent that decides to independently choose between the RL policy or a physics-based/rule-based simulator policy based on the given state information, previous experience and situation being modelled. The approach is similar to that demonstrated before in the robotics and transfer learning domains as explained in [25] and [27] where the problem was much simpler due to being a smaller domain. A confidence measure is defined and associated with the RL policy. The agent chooses to follow the RL policy if it is confident of the policy due to sufficient exploration on similar scenarios. Otherwise, it follows the physics based policy and uses these actions to refine the RL policy. This approach helps the RL policy to be improved faster by targeted exploration instead of random exploration. Additionally, the execution time is reduced and accuracy improved due to less reliance on the physics based policy.

4 Conclusions

In our paper, we have put forth a novel approach for utilizing RL in learning wildfire spread dynamics. The most important advantage of our method against the contemporary fire simulators is that it does not suffer from the errors due to biases in expert rules towards fires in North America as elaborated in [7]. Our approach is universal and can be used in any region given availability of sufficient data. As the fire policy is learned by the agent, most uncertainties as described in [26] are also avoided. The intersection between the decision making tools of Artificial Intelligence, the pattern recognition tools of Machine Learning and the challenging datasets of sustainability domains offer a rich area for research. For the machine learning community our approach opens up new sets of challenging and plentiful data sets for learning patterns of spatial change over time in the form of spatially spreading wildfires and a platform for experimenting with new Deep RL approaches on a challenging problem with high social impact. For the wild fire researchers and practitioners, our approach would give a more robust and accurate system for wild fire prediction. We also hope that our work can lead to development of a comprehensive way of integrating Deep Learning and RL approaches to support the tasks of prediction, dynamics model learning and decision making in many different problems with SSP structure like floods, urban sprawl and disease spread studies.

References

[1] US Fire Administration. Fire statistics. `https://www.usfa.fema.gov/data/statistics/`. Accessed: 2018-01-01.

[2] Alan A Ager, Nicole M Vaillant, and Mark A Finney. Integrating fire behavior models and geospatial analysis for wildland fire risk assessment and fuel management planning. *Journal of Combustion*, 2011, 2011.

[3] The Atlantic. Why the wildfires of 2018 have been so ferocious. `https://www.theatlantic.com/science/archive/2018/08/why-this-years-wildfires-have-been-so-ferocious/567215/`. Accessed: 2018-09-25.

[4] Sudipto Banerjee, Alan E Gelfand, Andrew O Finley, and Huiyan Sang. Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4):825–848, 2008.

[5] Paulo Cortez and Anibal de Jesus Raimundo Morais. A data mining approach to predict forest fires using meteorological data. In *volume proceedings of the 13th Portuguese conference on Artificial Intelligence*, 2007.

[6] Miguel Cruz and Martin Alexander. Assessing crown fire potential in coniferous forests of western north america: A critique of current approaches and recent simulation studies. *The Bark Beetles, Fuels, and Fire Bibliography*, 19, 01 2010.

[7] Miguel G Cruz and Martin E Alexander. Assessing crown fire potential in coniferous forests of western north america: a critique of current approaches and recent simulation studies. *International Journal of Wildland Fire*, 19(4):377–398, 2010.

[8] Mark A Finney. Farsite: Fire area simulator-model development and evaluation. *Res. Pap. RMRS-RP-4, Revised 2004. Ogden, UT: US Department of Agriculture, Forest Service, Rocky Mountain Research Station. 47 p.*, 4, 1998.

[9] N Forsell, F Garcia, and R Sabbadin. Reinforcement learning for spatial processes. In *18th World IMACS/MODSIM Congress*, pages 755–761, 2009.

[10] Sriram Ganapathi Subramanian and Mark Crowley. Learning forest wildfire dynamics from satellite images using reinforcement learning. In *Conference on Reinforcement Learning and Decision Making*, Ann Arbor, MI, USA., 2017.

[11] Roghayeh Jahdi, Michele Salis, Ali A Darvishsefat, Fermin Alcasena, Mir A Mostafavi, V Etemad, Olga M Lozano, and Donatella Spano. Evaluating fire modelling systems in recent wildfires of the golestan national park, iran. *Forestry*, 89(2):136–149, 2015.

[12] Jon E Keeley. Fire intensity, fire severity and burn severity: a brief review and suggested usage. *International Journal of Wildland Fire*, 18(1):116–126, 2009.

[13] Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pages 282–293. Springer, 2006.

[14] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937, 2016.

[15] Global News. Here's why b.c.'s 2018 wildfire season is so extreme, and why it's not over yet. `https://globalnews.ca/news/4390372/bc-wildfire-extreme-not-over-yet/`. Accessed: 2018-09-25.

[16] George D Papadopoulos and Fotini-Niovi Pavlidou. A comparative review on wildfire simulators. *IEEE systems Journal*, 5(2):233–243, 2011.

[17] Marc-André Parisien, KG Hirsch, SG Lavoie, JB Todd, VG Kafka, et al. Saskatchewan fire regime analysis. *Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre Information Report NOR-X-394.(Edmonton, AB)*, 2004.

[18] Marc-André Parisien, VG Kafka, KG Hirsch, JB Todd, SG Lavoie, PD Maczek, et al. Mapping wildfire susceptibility with the burn-p3 simulation model. *Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre Edmonton (AB)*, 2005.

[19] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006.

[20] Mahsa Salehi, Laura Irina Rusu, Timothy Lynar, and Anna Phan. Dynamic and robust wildfire risk prediction system: an unsupervised approach. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 245–254. ACM, 2016.

[21] Sriram Ganapathi Subramanian and Mark Crowley. Combining mcts and a3c for prediction of spatially spreading processes in forest wildfire settings. In *Canadian Conference on Artificial Intelligence*, Toronto, Ontario, Canada, 2018. Springer.

[22] Sriram Ganapathi Subramanian and Mark Crowley. Using spatial reinforcement learning to build forest wildfire dynamics models from satellite images. *Frontiers in ICT: Environmental Informatics*, 2018.

[23] Richard S Sutton. Introduction: The challenge of reinforcement learning. In *Reinforcement Learning*, pages 1–3. Springer, 1992.

[24] Richard S Sutton and Andrew G Barto. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.

[25] Matthew E Taylor, Halit Bener Suay, and Sonia Chernova. Integrating reinforcement learning with human demonstrations of varying ability. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 617–624. International Foundation for Autonomous Agents and Multiagent Systems, 2011.

[26] Matthew P Thompson and Dave E Calkin. Uncertainty and risk in wildland fire management: a review. *Journal of environmental management*, 92(8):1895–1909, 2011.

[27] Zhaodong Wang and Matthew E Taylor. Improving reinforcement learning with confidence-based demonstrations. In *Proceedings of the 26th International Conference on Artificial Intelligence (IJCAI)*, 2017.