

Physics-Informed Machine Learning Simulator for Wildfire Propagation

Luca Bottero¹, Francesco Calisto¹, Giovanni Graziano¹, Valerio Pagliarino¹,
Martina Scauda², Sara Tiengo¹, Simone Azeglio¹

¹Student, Department of Physics, University of Turin

²Student, Department of Mathematics, University of Turin

luca.bottero192, francesco.calisto, giovanni.graziano136, valerio.pagliarino, martina.scauda, sara.tiengo,
simone.azeglio@edu.unito.it

Abstract

The aim of this work is to evaluate the feasibility of re-implementing some key parts of the widely used Weather Research and Forecasting WRF-SFIRE simulator by replacing its core differential equations numerical solvers with state-of-the-art physics-informed machine learning techniques to solve ODEs and PDEs, in order to transform it into a real-time simulator for wildfire spread prediction. The main programming language used is Julia, a compiled language which offers better performance than interpreted ones, providing Just in Time (JIT) compilation with different optimization levels (Lubin and Dunning 2015). Moreover, Julia is particularly well suited for numerical computation and for the solution of complex physical models, both considering the syntax and the presence of specific libraries such as *DifferentialEquations.jl* and *ModelingToolkit.jl*.

INTRODUCTION

In recent years wildfires have been increasingly growing in intensity and frequency, becoming a serious threat to the health and socio-economic stability of various countries all over the world. In particular, Australia was devastated by the "Black Summer", the bushfire season between 2019 and 2020, and California just suffered from the most severe wildfire season recorded in its modern history. According to the California Department of Forestry and Fire Protection (Gauk-Roger et al. 2020): over 4 percent of its land was burned by more than 8,600 fires (California 2020).

The interconnection between wildfires and climate change is evident and has been studied by climate scientists over the years (Turco et al. 2014; Romps et al. 2014). On one hand, climate change favors the spread of wildfires. One of the main factors in play is the rise in global temperatures: the early beginning of spring leads to the rapid melting of snowpacks, causing land to dry out earlier and remain dry for longer. In addition, the hotter the air, the more water it soaks up from plants and soils; the "vapor pressure deficit" is used to measure the difference between how much water the air holds and how much it could hold: the higher this coefficient, the more soil and vegetation will dry out (Borunda 2020). Furthermore, bark beetles and other

insects that survive at high temperatures are responsible for killing millions of trees, according to the Fourth National Climate Assessment (Reidmiller et al. 2018), turning them into kindling for wildfires. On the other hand, wildfires impact climate change in two ways: firstly, they emit massive amounts of carbon dioxide and other pollutants that can affect regional and even global climate. Secondly, vast lands laid bare by fires are soaked with rain, increasing the risk for devastating landslides. On top of this, wildfires produce serious health risks due to the emission of toxic chemical compounds (Lieberman 2019).

In such a worrisome scenario, having the support of fast and accurate wildfire spread simulators that can run **iteratively** in order to evaluate several different containment strategies is of unprecedented importance.

Real-time is revolutionary

The current stable version of WRF is able to produce highly accurate and validated predictions, but each simulation takes several hours, which is acceptable for a single run. Nevertheless, if one wants to use the simulator for containment strategies, time is a major obstacle. In order to use WRF for that purpose, one should run the simulator iteratively with different initial scenarios, therefore predicting the outcome of various containment strategies and choosing the best one according to the output of the program. The revolutionary reach of such an approach is that it would limit human error in situations where taking the right decision quickly is extremely difficult. To reach this goal it is essential to significantly increase the computational efficiency of fire spread models, which can be done thanks to Machine Learning techniques.

Physics-based or ML?

In the context of science, the well-known adage "a picture is worth a thousand words" might well be "a model is worth a thousand datasets.", writes Christopher Rackauckas in his trademark paper (Rackauckas et al. 2020). A single sentence couldn't describe the true power of Scientific Machine Learning better. Machine Learning's utmost strength, i.e. universal flexibility to approximate any nonlinearity from data, as stated by the Universal Approximation Theorem (Kůrková 1992), is at the same time its main drawback: the

higher the complexity of the problem, the more data the algorithm needs to train. While some fields can provide big data, such as bioinformatics (Li et al. 2019), many scientific disciplines are still limited in the use of Deep Learning (DL) by a lack of data. Another reason for the ineffectiveness of Machine Learning (ML) might be the apparent chaotic nature of a phenomenon. This is the case of wildfires and weather-related phenomena, which don't depend solely on the initial and boundary conditions, but evolve according to complex internal dynamics that cannot be predicted from mere data. In order to make a neural network learn such dynamics, one would need data with both high spatial and temporal resolution, something not available in this field (Cenci et al. 2018). For this reason, we think that the growing field of Scientific Machine Learning is a good compromise where climate and wildfire science could take advantage. The idea is to embed physical information, generally in the form of partial differential equations, into the ML model, thus constructing a Physics Informed Neural Network (PINN). A former physics-informed approach consists in training the model on data and constraining the solution space to the physically admissible ones (e.g. for incompressible fluid dynamics, any flow solution that breaks the mass conservation principle would not be considered by the model (Raissi, Perdikaris, and Karniadakis 2019)). Another physics-informed architecture - the one considered in our case - consists instead in training the neural network directly on the physical model output data, employing the feed-forward networks as an efficient approximator of trajectories in the solution space. This approach leads to various advantages, such as avoiding the so-called *curse of dimensionality*: as the dimensionality of an equation increases, the computational costs of a numerical equation solver grow exponentially, while PINNs can be proven to have a polynomial bound (Grohs, Jentzen, and Salimova 2019). This results in a significant speed up with respect to numerical solvers when solving high dimensional problems. We decided to use the second technique since our goal is to evaluate the feasibility of a high efficiency, ML-based, wildfire spread simulator that uses the equations of the WRF-SFIRE model.

METHODS

Firstly, we started by investigating the structure of the systems of partial differential equations (PDEs) that constitute the core of the physical model of WRF-SFIRE. In this way we detected the modules where replacing the current structure has the greatest impact (i.e. CPU overhead). As expected, numerical solvers of the model's main differential equations are responsible for more than 20% of the total overhead. In particular, the level set equation (a PDE which describes the spread of the fireline, according to the level set method) and the Euler system (a system of 7 PDEs governing the atmosphere behavior), have significant impacts on the CPU load during fire simulation. **Therefore, in this paper we will present a detailed application of the PINN architecture to the level set equation and we will outline some ideas about the implementation of Euler system.**

Mathematical structure

Atmospheric governing equations The flux form conserved prognostic variables are defined as

$$\begin{aligned} \mathbf{V} &= \mathbf{d}\mathbf{v} = (U; V; W) \\ &= \mathbf{d}! \\ Q_m &= \mathbf{d}q_m \end{aligned} \quad (1)$$

where $\mathbf{d} = \partial/\partial p_d$ is the vertical coordinate metric, $\mathbf{v} = (U; V; W)$ are the covariant velocities in the horizontal and vertical directions, $! = \partial/\partial p$ is the contravariant vertical velocity, q_m is the moist potential temperature and $q_v; q_c; q_r; \dots$ denote the mixing ratios of moisture variables. Accordingly, the flux-form Euler equations, representing the core of the ARW, can be written as follows (Skamarock et al. 2019):

Conservation of momentum

$$\begin{aligned} \partial_t U + (r - V)U + \mathbf{d}_x p + \left(\frac{r}{\mathbf{d}}\right) \partial_x p @ = F_U \\ \partial_t V + (r - V)V + \mathbf{d}_y p + \left(\frac{r}{\mathbf{d}}\right) \partial_y p @ = F_V \\ \partial_t W + (r - V)W - g \left[\left(\frac{r}{\mathbf{d}}\right) \partial p - \mathbf{d}\right] = F_W \end{aligned} \quad (2)$$

Conservation of heat

$$\partial_t m + (r - V)m = F_m \quad (3)$$

Conservation of mass

$$\partial_t d + (r - V)d = 0 \quad (4)$$

Geopotential Material Derivative

$$\partial_t + \mathbf{d}^1 [(V - r) - gW] = 0 \quad (5)$$

Scalar moisture equations

$$\partial_t Q_m + (r - V)Q_m = F_{Q_m} \quad (6)$$

where

$$r - Va = \partial_x(Ua) + \partial_y(Va) + \partial_z(a)$$

and

$$V - ra = U\partial_x a + V\partial_y a + \partial_z a;$$

for a generic variable a .

These equations are coupled with the *diagnostic equation for dry hydrostatic pressure*

$$\partial = \mathbf{d}d \quad (7)$$

and the *diagnostic relation for the full pressure* (dry air combined with water vapor)

$$p = p_0 \frac{R_d m}{p_0 d} \quad (8)$$

In Eq. (2) - (6), $\mathbf{d} = \frac{1}{p_d}$ is the inverse density of dry air, $\mathbf{d} = \mathbf{d}(1 + q_v + q_c + q_r + \dots)^{-1}$, while R_d is the gas constant for dry air and $\frac{c_p}{c_v} = 1.4$ is the ratio of the heat capacities for dry air. The right hand side terms $F_U; F_V; F_W; F_m$ contain the Coriolis and curvature terms along with mixing terms and physical forcings.

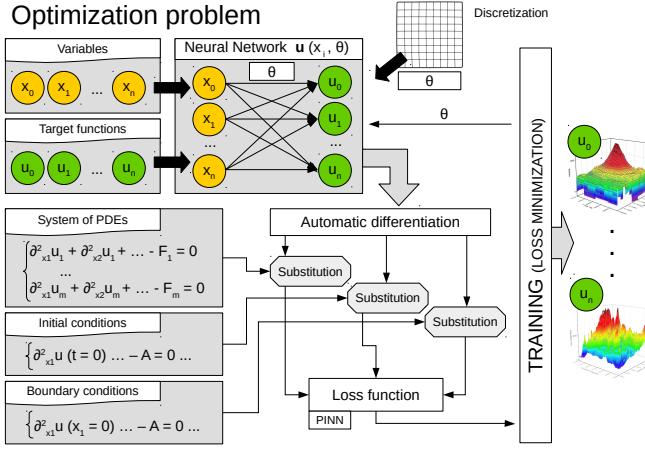


Figure 1: Physics-informed neural network architecture for solving PDEs

Level set equation In SFIRE (Mandel, Beezley, and Kochanski 2011), the propagation of a fire burning in the area $\Omega = \Omega(t)$ in the horizontal $(x; y)$ plane on which the Earth is projected, is implemented by the level set method, which evolves a function $\phi = \phi(x; t)$, called the *level set function*, such that the burning area at a time t is

$$\Omega(t) = \{x : \phi(x; t) > 0\}$$

and the fireline $\Gamma(t)$, i.e. the boundary of the burning region $\Omega(t)$, is the level set

$$\Gamma(t) = \{x : \phi(x; t) = 0\} \quad (9)$$

The evolution of the level set function is governed by the partial differential equation

$$\phi_t + S|\nabla\phi| - \Delta\phi = 0; \quad (10)$$

called the *level set equation*. The model adopts a semi-empirical approach, where the fire spread rate S is computed from fuel properties, using the following modified Rothermel formula (Rothermel 1972):

$$S = R_0(1 + \alpha W + \beta S) \quad (11)$$

Here, R_0 represents the spread rate in the absence of wind, whereas αW and βS are respectively the wind factor and the slope factor. In WRF, this equation is solved numerically through discretization on the refined fire grid, adopting Heun's method, a second-order Runge Kutta method ((Mandel, Beezley, and Kochanski 2011), p.596).

The new architecture

We have replaced the original numerical solvers with a profoundly different approach, by taking advantage of the recently developed Julia packages *NeuralPDE.jl* and *DiffEqFlux.jl* (Rackauckas et al. 2019), which support Physics Informed Neural Networks (PINNs) for automated PDE solving and Backwards Stochastic Differential Equation (BSDE)

methods to deal with parabolic PDEs. We employed the former method: PINNs are feed forward neural networks that are trained to solve supervised learning tasks while respecting any given law of physics described by general nonlinear partial differential equations. The resulting neural networks form a new class of universal function approximators that naturally encode any underlying physical laws as prior information (Zubov 2020).

The underlying architecture can be summarized by the following steps, with reference to Fig.1. Let's put ourselves in the most general case, where we want to find n target functions $[u_1; \dots; u_n]$ which satisfy a system of n PDEs. We first construct a surrogate solution to $u(x)$ as a neural network $\hat{u}(x; \theta)$ with n inputs and outputs and parameters θ ; since a neural network is mathematically a composite function, the derivatives of \hat{u} with respect to its input can be evaluated by applying the chain rule for differentiating compositions of functions using automatic differentiation (AD). Then we need to define a loss function (we used the L^2 norm) in order to calculate the discrepancy between $[u_1; \dots; u_n]$ and the outputs of the neural network: this is where we force our network to satisfy the physics imposed by the PDEs. Indeed, the loss is defined by substituting our neural network and its derivatives back into the equations we want to solve (where we brought all the terms on the left-hand side of the equal sign). The last step is to train the neural network to find the best parameters by minimizing the loss function with gradient-based optimizers, such as Adam or LBFSGS. It is straightforward that minimizing the loss function means solving the equations, since we are approximating our target functions with \hat{u} better and better. **The key point is that it converts an integration problem into a mere minimization task and it does not need data to train, because it is trained on the PDEs themselves.**

Implementation of the Level set

The level-set differential equation is, as previously explained, the mathematical core for calculating the spread of the fire. The level set equation (that is a 2D surface embedded in a 3D euclidean space) is initially set as the distance of each point from the ignition's line. Our approach requires that every quantity is provided either as a constant or as a function. Therefore we implemented the initial condition as a cone with elliptical section, because we had to deal with circular and linear ignition shapes (theoretically, any ignition shape can be given as input). The cone is then shifted down in order to make the contour level $z = 0$ resembling the desired initial fire line (see the Appendix for a graphical representation). Multiple ignition points are possible, but further refinements are needed to make this possibility fully functional. The algorithm contains all the necessary expressions for the calculation of the fire spread rate, taking into account the values of the wind, the altimetric profile, the fuel map etc. These variables have to be given in input as differentiable vector fields of space and time. This can be considered both a limit and an advantage, in fact on one hand the discretized measurements must be fitted using approximating functions, but on the other hand this method is very computationally efficient. The possibility to give in

Optimization Algorithm	ADAM (4800 iterations)
Final Objective Value	6.39 e-8
Training Strategy	QuadratureTraining()
Domains	$t \geq [0;10], x \geq [0;10], y \geq [0;10]$
Training Mesh Size	$[dt; dx; dy] = [0.17; 0.02; 0.02]$
Initial Condition	$(0; x; y) = (5(x - 0.3)^2 + 0.15y^2)^{\frac{1}{2}}$ 0.2
Neural Network dimensions	$3 > 16 > 1$
Training Time	647 s

Table 1: Technical Specs of the model - One Fire
Domain sizes are expressed in arbitrary units, since scaling factors transform input and output in pre-processing and post-processing phases.

input finite matrix of values is presented in "Future Work" and depends on the fact that some key libraries we use are still under development. You can see in the appendix an example of the altimetric profile of the *Isom Creek* fire location fitted using a 4th degree polynomial. The last step before the training of the model is a proportional scaling of the values to a domain with size in the order of 10^1 and the definition of a discretization that is used only to run the training algorithm. The outputs will be continuous functions. The discretization to this interval prevents some numerical instabilities and spikes that are linked to the early stage of development of some Julia libraries in use. The loss function is made of two parts: the former minimizing the L^2 error referred to the PDE definition and the latter minimizing the L^2 distance between the boundary conditions and the solution evaluated at the boundaries. At this point the training algorithm is instantiated and launched. The minimization of the loss functions is the process that actually solves the PDE and constitutes the main load for the CPU. It can be easily parallelized and thus accelerated using GPUs. When the training is completed the prediction undergoes a new proportional scaling that gives back the original domain shape.

Ideal simulation: "One fire"

For clarity, we renamed the *Two Fires* ideal simulation as *One Fire*, after having removed one ignition point. It is a simple case that demonstrates that our model works well on the level set in case of fires that spread in flat areas where the fuel doesn't change. It also evolves according to the wind direction, represented by a vector with two components. The results at different time frames are represented in Fig.2. The details about the architecture employed and the training are listed in TABLE 1.

In order to provide a quantitative measure of the error between the outputs we decided to use the Hausdorff distance

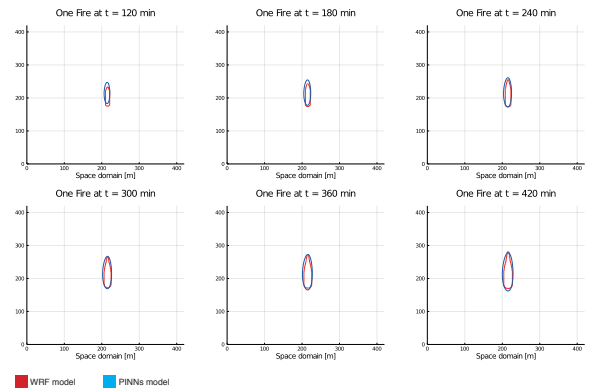


Figure 2: Comparison between the WRF and PINNs outputs at different time frames




Figure 3: Hausdorff distance as a measure of the error - One Fire.

(Knauer et al. 2009),

$$d_H(X; Y) = \max \left(\sup_{x \in X} \inf_{y \in Y} d(x; y); \sup_{y \in Y} \inf_{x \in X} d(x; y) \right);$$

normalized on the area of the curves that represent the fire-line, as Fig.3 shows. The fact that the error decreases tells us that as the area grows with time, the difference between the two regions doesn't increase, meaning that they evolve in synchrony and with the same shape.

Simulation of a real wildfire: Isom Creek, Alaska, 6/5/2020

We have been able to simulate the Isom Creek fire (Incident-Information-System June 2020) both with WRF and with our implementation. Although it is not possible to perform a precise quantitative comparison between WRF and the real data, or between our result and the real data (due to a lack of temporal precision of the available data at (Incident-Information-System 2020)), it is clear that the WRF output simulates the wildfire accurately, at least visually. The result obtained with the PINN architecture is shown in Fig.4. It is clear that our implementation hasn't been able to capture the shape of the fireline with the same accuracy. This will be pointed out below as a temporary limitation to our

Optimization Algorithm	ADAM (3000 iterations)
Final Objective Value	2.14 e-6
Training Strategy	QuadratureTraining()
Domains	$t \in [0;10], x \in [0;10], y \in [0;10]$
Training Mesh Size	$[dt; dx; dy] = [0.17; 0.02; 0.02]$
Initial Condition	$(0; x; y) = ((x - 5)^2 + 0.7(y - 5)^2)^{\frac{1}{2}}$ 0:2
Neural Network dimensions	$3 > 16 > 1$
Training Time	270 s

Table 2: Technical Specs of the model - The considerations on measurement units made in TABLE 1 apply also here.

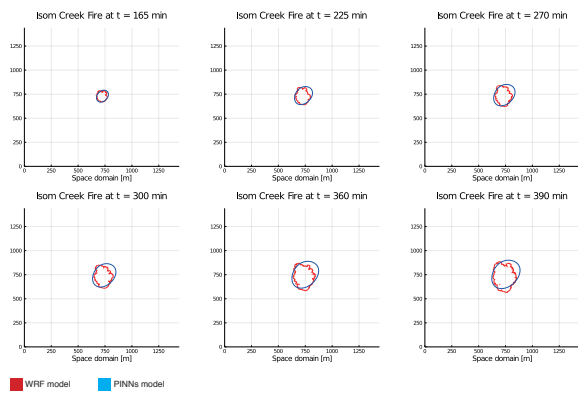


Figure 4: Comparison between the WRF and PINNs outputs at different time frames

model: at the moment it is not possible to assign a different fuel category to each mesh point (i.e. using a matrix to build a function), so the fireline will always have a smooth perimeter. The specifics of the architecture are outlined in TABLE 2.

The error can be quantitatively measured with the Hausdorff distance as before, with the outcome plotted in Fig.5.

The considerations outlined in the previous paragraph remain valid here.

DISCUSSION

Implementation of the Euler system

On top of the results previously outlined, we have also taken on the implementation of the most computationally expensive part of the WRF physical model. Writing and solving the 7-equation Euler system in Julia was really challenging, in fact at the moment we are not aware of any publications where these techniques are yet applied to PDE systems of such complexity. The first problem was to figure out which were the independent and dependent variables: climate models adopt different conventions depending

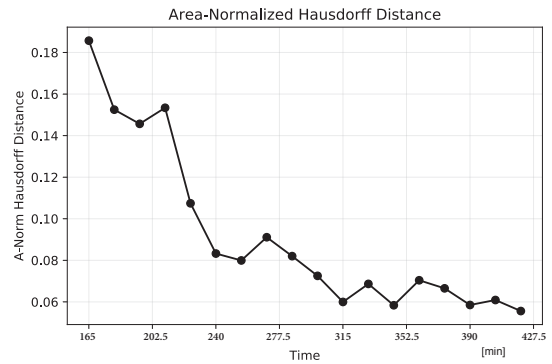


Figure 5: Hausdorff distance as a measure of the error - Isom Creek.

on the circumstances (e.g., they generally use pressure coordinates instead of height coordinates). Then we had to choose the most suitable training strategy, which turned out to be *QuadratureTraining*. We have used parabolic initial conditions for a first evaluation, but we want to investigate them more in depth, since the study of boundary conditions often deserves a paper on its own. Unfortunately, the *NeuralPDE* library is still unable to treat this kind of problem with stability, and often incurs errors due to internal divergence calculations. Despite this, we have been able to obtain convergence of the loss function, although it is not enough to present valid results. We contacted the authors of these libraries, that are still under development, and we are looking forward to contributing.

The results exposed above show the feasibility of utilizing the Julia *NeuralPDE* library to model a wildfire related differential equation, i.e. the level set equation. As disclosed above, we have also obtained some preliminary but unstable results regarding the atmospheric coupling of the model by attempting to solve the Flux-Form Euler system equation. The outputs obtained with our approach don't strictly follow the jagged firelines of WRF: that's not an intrinsic limit of the architecture, which doesn't lack variance even if it is rather simple; the reason is, as explained above, that the inputs (fuel and wind in particular) we gave to the NN are not varied enough, due to limitations of the library. In spite of this we have shown the ability of our implementation to reproduce the desired time evolution of the fireline as given by WRF.

Advantages of this architecture

The various advantages of this approach are:

- We can get a rough estimate of the speed up our model provides - compared to numerical solvers - by noting that the train time is of some hundreds of seconds, while the WRF run time is more than an hour. This point has to be further studied: on one hand, we used a simpler set of inputs for the reasons we will explain later; on the other hand we worked on one of the eight cores at our disposal.
- While numerical methods only solve on the initial domain

range, this architecture learns the neural network parameters that can be used to **extend the predicted solution outside of the training domains** (with a decrease in accuracy if the extrapolation is extended too much). This can be used for "forensic investigations": it is possible to reconstruct the ignition point and the evolution of the fire-line from the final state of the wildfire.

- It is possible to interpolate the solution on a continuous mesh.
- Modifying the equations of the model is as easy as changing a few lines of code, instead of re-implementing the discretization and perturbative form of the equations, allowing for faster investigation of new models.
- Given the possibility to use several CPU cores, the speed could increase even more and this approach could be run iteratively in order to simulate the outcome of different *containment scenarios* and choose the best among them.
- This architecture doesn't need neural networks with a high number of degrees of freedom to be accurate, reducing the overall computational cost and increasing performance. In quantitative terms, this means not exceeding some tens of neurons on a single hidden layer.
- As a consequence, the model is more interpretable (it is less similar to a complete *black-box*) and its users can run the resulting simulator on standard PCs without needing high performance machines.
- PINNs can be used to extrapolate physical laws once they have been interpolated, hence providing a tool to refine theoretical models. However, this requires a greater amount of input data.

Limitations

Let's now analyze the limits of our approach:

- At the moment, it cannot run on GPUs nor multiple CPUs because the library we used is not fully developed yet. We contacted the developers and they are working on it.
- The library doesn't support the usage of matrices as functions of their indices. This means it is not possible to use a variable fuel map, fundamental in order to capture the evolution of the fire with higher accuracy. We opened an issue about this (Calisto 2020).
- Our approach suffers from numerical instability depending on the training strategy adopted and on the number of neurons and layers.
- The model only works on relatively small domains and is unstable for larger ones. Since the output has a small codomain, we had to scale the quantities involved up and introduce a scale factor on the fire spread rate in order to make it compatible with the evolution of the output of WRF.
- We still have to figure out how to predict the spread of a wildfire with more than one ignition point.
- The architecture has not yet been extended to Convolutional Neural Networks, Recurrent Neural Networks or others.

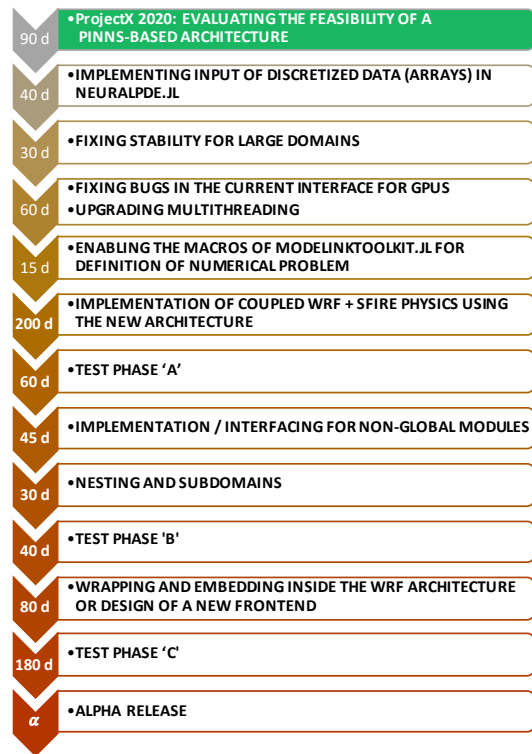


Figure 6: Future prospect of our work

FUTURE WORK

This research is the first step towards a concrete implementation of PINNs-based wildfire simulation. We have designed a chart (Fig.6) that shows the next improvements to be made in order to achieve such a ambitious goal.

CONCLUSIONS

The study we carried out had the goal to investigate the applicability of the recently developed field of Scientific Machine Learning on climate, wildfires in particular, models. We have outlined some results that tell us that many improvements are needed in order to transform this into a validated product, but also show the big potential of our approach. We need to add further refinements to the implementation in order to carry out a precise time comparison between the two approaches, but the results obtained thus far show promising evidence. The encouraging outcome inspires us to continue our work by improving the architectures and possibly employ them in different fields of research. We hope that this line of research will be considered as a starting point for a more effective cohesiveness between Machine Learning and Physical Models in Climate Science, and thus further explored by other researchers.

Acknowledgments

This work was presented at the ProjectX 2020 competition by UofT AI. We acknowledge University of Turin, Machine Learning Journal Club for supporting. We thank Professor Enrico Ferrero (Università del Piemonte Orientale),

Professor Massimiliano Manfrin (University of Turin) and the whole Atmospheric Physics and Meteorology Group, PhD Christopher Rackauckas (Massachusetts Institute of Technology), PhD Kirill Zubov (Saint-Petersburg State University), Vaibhav Dixit (Julia Computing), PhD Brian Wee (Founder at Massive Connections), Dr. Rustem Arif Albayrak (NASA), PhD David Marvin (CEO at Salo Sciences), Professor Piero Fariselli (University of Turin) and Pietro Monticone M.Sc. student (University of Turin), for their precious help and availability. We acknowledge the company *Mollificio Astigiano* (Belveglio, Asti, Italy) for providing the computational power needed for this research and the *HPC4AI* center of the University of Turin for their support.

References

- Borunda, A. 2020. The science connecting wildfires to climate change. URL <https://www.nationalgeographic.com/science/2020/09/climate-change-increases-risk-fires-western-us/>.
- California. 2020. Fire statistics, CALFIRE. URL <https://www.fire.ca.gov/incidents/2020/>.
- Calisto, F. 2020. Embedding arrays into NeuralPDE differential equations. URL <https://github.com/SciML/NeuralPDE.jl/issues/177>.
- Cenci, L.; Pulvirenti, L.; Boni, G.; and Pierdicca, N. 2018. Defining a Trade-off Between Spatial and Temporal Resolution of a Geosynchronous SAR Mission for Soil Moisture Monitoring. *Remote Sensing* 10: 1950. doi:10.3390/rs10121950.
- Gauk-Roger, T.; Chan, S.; Hanna, J.; and Almasry, S. 2020. California wildfires: Fire chief says dozens of major blazes have state in 'dire situation'. *CNN* URL <https://edition.cnn.com/2020/09/08/us/california-fires-tuesday/index.html>.
- Grohs, P.; Jentzen, A.; and Salimova, D. 2019. Deep neural network approximations for Monte Carlo algorithms.
- Incident-Information-System. 2020. Isom Creek Fire Progression Map. URL <https://inciweb.nwccg.gov/incident/map/6758/798095>.
- Incident-Information-System. June 2020. Isom Creek Fire Information. URL <https://inciweb.nwccg.gov/incident/6758>.
- Knauer, C.; Löffler, M.; Scherfenberg, M.; and Wolle, T. 2009. The directed Hausdorff distance between imprecise point sets.
- Kůrková, V. 1992. Kolmogorov's theorem and multilayer neural networks. *Neural Networks* 5. doi:10.1016/0893-6080(92)90012-8.
- Li, Y.; Huang, C.; Ding, L.; Li, Z.; Pan, Y.; and Gao, X. 2019. Deep learning in bioinformatics: Introduction, application, and perspective in the big data era. *Methods* 166: 4–21.
- Lieberman, B. 2019. Wildfires and climate change: What's the connection? URL <https://yaleclimateconnections.org/2019/07/wildfires-and-climate-change-whats-the-connection>.
- Lubin, M.; and Dunning, I. 2015. Computing in Operations Research Using Julia. *INFORMS Journal on Computing* 27(2): 238–248. doi:10.1287/ijoc.2014.0623.
- Mandel, J.; Beezley, J. D.; and Kochanski, A. K. 2011. Coupled atmosphere-wildland fire modeling with WRF 3.3 and SFIRE 2011. *Geoscientific Model Development* 4(3): 591–610. doi:10.5194/gmd-4-591-2011.
- Rackauckas, C.; Innes, M.; Ma, Y.; Bettencourt, J.; White, L.; and Dixit, V. 2019. DiffEqFlux.jl - A Julia Library for Neural Differential Equations. *CoRR* abs/1902.02376.
- Rackauckas, C.; Ma, Y.; Martensen, J.; Warner, C.; Zubov, K.; Supekar, R.; Skinner, D.; and Ramadhan, A. 2020. Universal differential equations for scientific machine learning. *arXiv preprint: 2001.04385*.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378: 686–707.
- Reidmiller, D. R.; Avery, C. W.; Easterling, D. R.; Kunkel, K. E.; Lewis, K. L.; Maycock, T. K.; and Stewart, B. C. 2018. Impacts, Risks, and Adaptation in the United States: The Fourth National Climate Assessment, Volume II. Technical report. doi:10.7930/nca4.2018.
- Romps, D. M.; Seeley, J. T.; Vollaro, D.; and Molinari, J. 2014. Projected increase in lightning strikes in the United States due to global warming. *Science* 346(6211): 851–854. ISSN 0036-8075. doi:10.1126/science.1259100. URL <https://science.sciencemag.org/content/346/6211/851>.
- Rothermel, R. C. 1972. A Mathematical Model for Predicting Fire Spread in Wildland Fires. *USDA Forest Service Research Paper INT-115* URL <http://www.treesearch.fs.fed.us/pubs/32533>.
- Skamarock, W. C.; Klemp, J. B.; Dudhia, J.; Gill, D. O.; Liu, Z.; Berner, J.; Wang, W.; Powers, J. G.; Duda, M. G.; Barker, D. M.; and Huang, X.-Y. 2019. A Description of the Advanced Research WRF Model Version 4. doi:10.5065/1DFH-6P97.
- Turco, M.; Llasat, M.-C.; von Hardenberg, J.; and Provenzale, A. 2014. Climate change impacts on wildfires in a Mediterranean environment. *Climatic Change* 125(3-4): 369–380. doi:10.1007/s10584-014-1183-3.
- Zubov, K. 2020. Physics-informed neural networks (PINNs) solver on Julia. GSoC 2020. First evaluation. URL https://nextjournal.com/kirill_zubov/physics-informed-neural-networks-pinns-solvers-on-julia-gsoc-2020-first-evaluations.

APPENDIX

All the code produced is published on a public repository available at this link: <https://github.com/MachineLearningJournalClub/MLJC-UniTo-ProjectX-2020-public>. The output of the WRF simulations can be found in the following Google Drive folder: <https://tinyurl.com/mljc-unito-px2020>.

