Fusion of Physics-Based Wildfire Spread Models with Satellite Data using Generative Algorithms

Bryan Shaddy[1], **Deep Ray**[2], **Angel Farguell**[3], **Valentina Calaza**[1], **Jan Mandel**[4], **James Haley**[5], **Kyle Hilburn**[5], **Derek V. Mallia**[6], **Adam Kochanski**[3], **Assad Oberai**[1]
[1]Department of Aerospace and Mechanical Engineering, University of Southern California
[2]Department of Mathematics, University of Maryland, College Park
[3]Department of Meteorology and Climate Science, San Jose State University
[4]Department of Mathematical and Statistical Sciences, University of Colorado Denver
[5]Cooperative Institute for Research in the Atmosphere, Colorado State University
[6]Department of Atmospheric Sciences, University of Utah
Corresponding author: bshaddy@usc.edu

Abstract

Climate change has driven increases in wildfire prevalence, prompting development of wildfire spread models. Advancements in the use of satellites to detect fire locations provides opportunity to enhance fire spread forecasts from numerical models via data assimilation. In this work, a method is developed to infer the history of a wildfire from satellite measurements using a conditional Wasserstein Generative Adversarial Network (cWGAN), providing the information necessary to initialize coupled atmosphere-wildfire models in a physics-informed approach based on measurements. The cWGAN, trained with solutions from WRF-SFIRE, produces samples of fire arrival times (fire history) from the conditional distribution of arrival times given satellite measurements, and allows for assessment of prediction uncertainty. The method is tested on four California wildfires and predictions are compared against measured fire perimeters and reported ignition times. An average Sorensen's coefficient of 0.81 for the fire perimeters and an average ignition time error of 32 minutes suggests that the method is highly accurate.

1 Introduction

Climate change has induced increases in wildfire frequency and severity, with conditions conducive to large fires expected to become more common as climate predictions indicate drier summers and more heat in years to come [1–4]. Wildfires impact air quality, cause destruction of property and harm to health, and negatively influence atmospheric composition through the release of greenhouse gases, making wildfire behavior models critical [5–8]. Wildfires are strongly coupled to meteorological conditions, motivating the development of coupled wildfire-atmosphere models which aim to capture these interactions. However, issues of accumulated model errors leading to degraded forecasts exist, leading to interest in performing data assimilation [9–11]. Satellite-based active fire (AF) products are a commonly used measurement type for data assimilation which can provide AF detections with a resolution of up to 375 m at a temporal frequency of around 12 hours, however these measurements can have artifacts and are at spatio-temporal resolutions lower than current wildfire models [12].

For dynamic coupled atmosphere-wildfire spread models, the state of the system is represented by all the wildfire and atmosphere variables in the simulation. To perform data assimilation in this setting it has been demonstrated that if the precise history of a wildfire during its initial spread was known, then this history could be prescribed within a coupled wildfire-weather model to spin-up the atmosphere with the right amount of heat and mass flux added at the right place and time, yielding an atmospheric

Tackling Climate Change with Machine Learning workshop at NeurIPS 2023

state which is in sync with the fire state at the end of the initial phase, providing an initial condition to start a simulation from [13]. Thus, the data assimilation problem can be transformed to one of using satellite measurements to determine the history of a fire in the initial period, as succinctly represented by the fire arrival time which is defined as the time the fire arrives at a given location [13, 14].

Farguell et al. [15] used a support vector machine (SVM) to estimate fire arrival times based on satellite active fire data, however this method did not incorporate any physics inherent to wildfire spread into estimates and does not provide information about uncertainty. The method described here addresses these limitations, starting with a probabilistic interpretation of the problem where both the measured active fire pixels and the desired fire arrival time field are treated as random vectors. The inference problem is then one of quantifying the conditional probability distribution for the fire arrival time conditioned on a given measurement of active fire pixels. The measured fire pixels and the arrival time field are both high-dimensional random vectors, making this problem challenging to solve. To address this challenge, we utilize a conditional Wasserstein Generative Adversarial Network (cWGAN) to learn and then sample from a conditional probability distribution. The approach uses wildfire simulation results from WRF-SFIRE for training, allowing the appropriate physics to be injected into predictions. The trained algorithm is applied retrospectively to four California wildfires and its performance is assessed; additional comparisons are made to the SVM-based method.

2 Problem formulation

Let $\boldsymbol{\tau}$ denote the matrix of fire arrival times whose components τ_{ij} represent the fire arrival times for the i-th pixel along the longitude and the j-th pixel along the latitude. Assume there are N_τ such pixels and therefore $\boldsymbol{\tau} \in \Omega_{\boldsymbol{\tau}} \subset \mathbb{R}^{N_\tau}$. Each pixel is of size 60×60 meters. A measurement operator M may be applied which transforms $\boldsymbol{\tau}$ into a coarse, sparse, and noisy measurement $\bar{\boldsymbol{\tau}}$ using the mapping $M : \Omega_{\boldsymbol{\tau}} \to \Omega_{\bar{\boldsymbol{\tau}}}$. We note that M may easily be approximated, however here we are interested in the inverse problem, which maps from $\bar{\boldsymbol{\tau}}$ to $\boldsymbol{\tau}$ and is much more challenging to solve.

We recognize that a single measurement can correspond to a distribution of likely fire arrival times, and to cope with the ill-posed nature of this problem we adopt a probabilistic approach. We let the inferred field $\boldsymbol{\tau}$ and the measurements $\bar{\boldsymbol{\tau}}$ be modeled by random variables $\boldsymbol{T}$ and $\bar{\boldsymbol{T}}$, respectively. We recognize that given a measurement $\bar{\boldsymbol{\tau}}$, we are interested in learning and generating samples from the conditional distribution $P_{\boldsymbol{T}|\bar{\boldsymbol{T}}}$. We accomplish this by first generating N pairwise samples of arrival times and measurements $(\boldsymbol{\tau}^{(i)}, \bar{\boldsymbol{\tau}}^{(i)}),\ i = 1, \cdots, N$ from the joint distribution $P_{\boldsymbol{T}\bar{\boldsymbol{T}}}$ using WRF-SFIRE (and data augmentation) to generate N instances of arrival times $\boldsymbol{\tau}^{(i)}$, to which the measurement operator is applied to obtain the corresponding $\bar{\boldsymbol{\tau}}^{(i)}$. We then use this data to train a conditional Wasserstein Generative Adversarial Network (cWGAN). Following this, active fire satellite measurements of a wildfire may be used as input to the cWGAN to produce samples of the arrival time from the conditional distribution. These samples are then used to generate statistics of interest, including the pixel-wise mean and standard deviation in arrival time.

3 Conditional Wasserstein Generative Adversarial Networks (cWGAN)

The cWGAN consists of two subnetworks, a generator $\boldsymbol{g}$ and a critic d. The generator $\boldsymbol{g}$ is given by the mapping $\boldsymbol{g} : \Omega_{\boldsymbol{z}} \times \Omega_{\bar{\boldsymbol{\tau}}} \to \Omega_{\boldsymbol{\tau}}$, where $\boldsymbol{z} \in \Omega_{\boldsymbol{z}} \subset \mathbb{R}^{N_z}$ is a latent variable modeled using the random variable $\boldsymbol{Z}$, with distribution $P_{\boldsymbol{Z}}$ selected such that it is easy to sample from. The critic d is given by the mapping $d : \Omega_{\boldsymbol{\tau}} \times \Omega_{\bar{\boldsymbol{\tau}}} \to \mathbb{R}$.

For a given measurement $\bar{\boldsymbol{\tau}}$, the generator $\boldsymbol{g}$ produces samples $\boldsymbol{\tau}^{\boldsymbol{g}} = \boldsymbol{g}(\boldsymbol{z}, \bar{\boldsymbol{\tau}}),\ \boldsymbol{z} \sim P_{\boldsymbol{Z}}$ from the learned conditional distribution $P^{\boldsymbol{g}}_{\boldsymbol{T}|\bar{\boldsymbol{T}}}(\boldsymbol{\tau}|\bar{\boldsymbol{\tau}})$. The training of the cWAN requires this distribution to be close to the true conditional distribution $P_{\boldsymbol{T}|\bar{\boldsymbol{T}}}(\boldsymbol{\tau}|\bar{\boldsymbol{\tau}})$ in the Wasserstein-1 metric. The cWGAN is trained using the objective function, $\mathcal{L}(d, \boldsymbol{g}) = \mathbb{E}_{(\boldsymbol{\tau},\bar{\boldsymbol{\tau}})\sim P_{\boldsymbol{T}\bar{\boldsymbol{T}}}}[d(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}})] - \mathbb{E}_{\substack{\boldsymbol{\tau}^{\boldsymbol{g}}\sim P^{\boldsymbol{g}}_{\boldsymbol{T}|\bar{\boldsymbol{T}}} \\ \bar{\boldsymbol{\tau}}\sim P_{\bar{\boldsymbol{T}}}}}[d(\boldsymbol{\tau}^{\boldsymbol{g}}, \bar{\boldsymbol{\tau}})]$, and the optimal generator and critic ($\boldsymbol{g}^*$ and d^*, respectively) are determined by solving the min-max problem $(d^*, \boldsymbol{g}^*) = \arg\min_{\boldsymbol{g}}\ \arg\max_d\ \mathcal{L}(d, \boldsymbol{g})$. Assuming the critic d is 1-Lipschitz in both its arguments, it is shown in [16] that the $\boldsymbol{g}^*$ can be used to approximate the true conditional distribution. This implies that once the generator has been trained it may be used to approximate the statistics from

the true conditional distribution as $\mathbb{E}_{\boldsymbol{\tau}\sim P_{\boldsymbol{T}|\bar{\boldsymbol{T}}}}[\ell(\boldsymbol{\tau})] \approx \frac{1}{K}\sum_{i=1}^{K} \ell(\boldsymbol{g}^*(\boldsymbol{z}^{(i)}, \bar{\boldsymbol{\tau}})), \quad \boldsymbol{z}^{(i)} \sim P_{\boldsymbol{Z}}$. Following from this, the pixel-wise mean prediction for $\boldsymbol{\tau}$ based on a given $\bar{\boldsymbol{\tau}}$ can be computed by setting $\ell(\boldsymbol{\tau}) = \boldsymbol{\tau}$ and the pixel-wise variance may be computed by setting $\ell(\boldsymbol{\tau}) = (\boldsymbol{\tau} - \mathbb{E}[\boldsymbol{\tau}])^2$. The interested reader is referred to [17, 18, 16] for additional information. Further details of the cWGAN architecture are included in Appendix A.

3.1 Training the cWGAN

To train the cWGAN, samples of $\boldsymbol{\tau}$ are drawn from the prior marginal distribution $P_{\boldsymbol{T}}$ by performing simulations with the coupled atmosphere-wildfire model WRF-SFIRE. Fire arrival times are computed from 20 WRF-SFIRE simulations, each considering 2 day fire spread over flat terrain with uniform fuel and point ignitions. The simulations varied only in the initial wind speed, which was prescribed uniformly in one direction. Resulting $\boldsymbol{\tau}$ samples were on a grid of size 512×512, with a 60 m resolution. Data augmentation was performed by rotating and translating the 20 solutions $\boldsymbol{\tau}$. To each sample $\boldsymbol{\tau}$ an approximation of the measurement operator M was applied to generate a corresponding measurement $\bar{\boldsymbol{\tau}} = M(\boldsymbol{\tau})$, yielding the pairs $(\boldsymbol{\tau}^{(i)}, \bar{\boldsymbol{\tau}}^{(i)})$ from the joint distribution $P_{\boldsymbol{T}\bar{\boldsymbol{T}}}$ to be used for training. Lastly, a random amount of time was added to the fire arrival times to simulate an unknown ignition time. Additional details about the measurement operator are provided in Appendix B, along with sample data pairs in Appendix C. Training was tracked using a mismatch term defined as the 2-norm of the difference between the generated fire arrival time and the true fire arrival time. The cWGAN was trained in PyTorch for 200 epochs, using a batch size of 5 and the Adam optimizer.

4 Results

The cWGAN-based method is tested on four California wildfires, namely the Bobcat, Tennant, Oak, and Mineral fires. The method is applied to active fire measurements from the Visible Infrared Imaging Radiometer Suite (VIIRS) on board the polar-orbiting satellite Suomi-NPP, which provides detections with a resolutions of 375 m. Measurements were separated by confidence level, such that there is one measurement $\bar{\boldsymbol{\tau}}$ containing high confidence detections only and another measurement $\bar{\boldsymbol{\tau}}$ which contains both high and nominal confidence detections, resulting in two measurements $\bar{\boldsymbol{\tau}}$ per fire on which predictions may be conditioned. Specifics on the preprocessing of the 375 m Level-2 VIIRS AF data is provided in Appendix D, along with the measurements $\bar{\boldsymbol{\tau}}$ used here.

The two measurements for each fire are used as input to the trained cWGAN, generating 200 realizations of the arrival time per measurement by sampling the latent vector $\boldsymbol{z}$ from its distribution. It is important to note that predicted fire arrival times are relative to the start of the day. The realizations are combined with different weights (0.2 for the high confidence and 0.8 for the high + nominal confidence) to compute the pixel-wise mean and standard deviation of arrival times shown in Fig. 1. The mean arrival time may then be used to generate a smooth sequence of fire perimeters to initialize the state variables in a coupled weather/wildfire model, like WRF-SFIRE. It is additionally noted that the standard deviation plots in Fig. 1 provide a measure of uncertainty in the predictions. For comparison, Fig. 2 shows results from the SVM-based method described in Farguell et al. [15] for the same set of VIIRS AF data used with the cWGAN approach here. It is observed that the SVM method produces islands of unburnt regions within the predicted fire extents which are not present in the cWGAN predictions and further provides no measure of uncertainty for the predictions.

To quantitatively assess fire arrival time predictions, high resolution infrared (IR) wildfire extent perimeters provided by the National Infrared Operations (NIROPS) program are used as ground-truth, against which geolocated fire perimeters, produced by plotting a contour of the fire arrival time for a prescribed time, may be compared [19]. The IR perimeters used here are provided for the interested reader in Appendix D. Predicted perimeters are compared with measured perimeters by identifying the true positive pixels (burnt in both the prediction and truth), the false negative pixels (not burnt in the prediction but burnt in truth) and the false positive pixels (burnt in the prediction but not burnt in truth), marked as A, B and C, respectively, and using them to compute three metrics that quantify the prediction quality. These metrics are the Sørensen–Dice coefficient (SC), the Probability of Detection (POD), and the False Alarm Ratio (FAR), computed as: $SC = \frac{2A}{2A+B+C} \quad , \quad POD = \frac{A}{A+B} \quad , \quad FAR = \frac{C}{A+C}$. These metrics attain values between 0 and 1, where for the SC and POD the best model yields a value of 1 and for the FAR a value of 0 is ideal. Predicted ignition times,

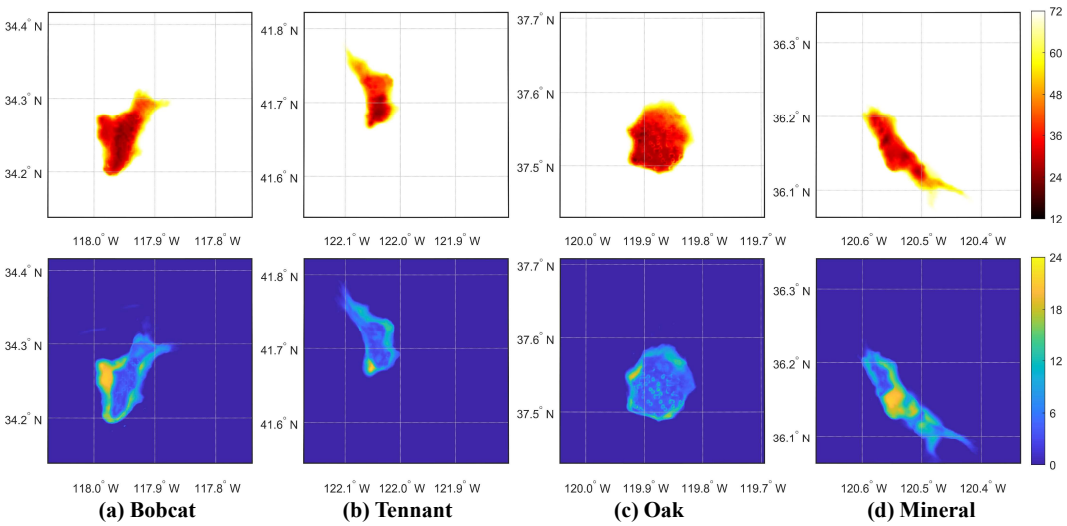

Figure 1: Weighted mean (first row) and standard deviation (second row) of fire arrival time predictions using the cWGAN approach.

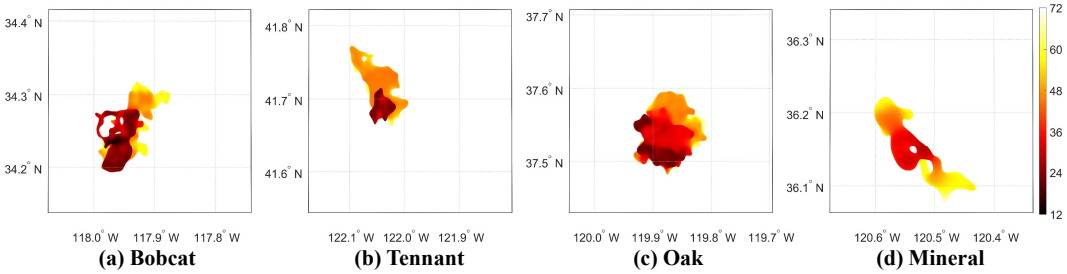

Figure 2: Fire arrival time predictions using the SVM-based approach described in Farguell et al. [15].

computed as the earliest predicted fire arrival time, may also be compared to ignition times reported by CAL FIRE, for fires where this information is available (Tennant, Oak, and Mineral). Table 1 contains the SC, POD, FAR, and ignition time error values computed for the Bobcat (SC, POD, and FAR only), Tennant, Oak, and Mineral fires based on predictions by the cWGAN method.

Quantitative results for the cWGAN predictions are compared to those for the SVM-based method, which are provided in Table 2. It is found that in all cases except the Tennant fire, the SC for the cWGAN method is higher than the SVM method. From the POD values we conclude that the SVM method performs better, indicating that the cWGAN is less likely than the SVM method to capture the full extent of the fire. On the other hand, the FAR values for the cWGAN are better than those for the SVM method in every case, indicating the cWGAN is less likely than the SVM method to suffer from false positive errors. Of these metrics, the SC is the most critical to do well on when evaluating

Table 1: Sørensen's coefficient (SC), probability of detection (POD), false alarm ratio (FAR), and ignition time error values obtained for the cWGAN predictions.

Wildfire	cWGAN SC	cWGAN POD	cWGAN FAR	cWGAN ignition time error
Bobcat	0.80	0.97	0.32	-
Tennant	0.78	0.78	0.21	41 minutes
Oak	0.84	0.97	0.26	20 minutes
Mineral	0.81	0.76	0.14	36 minutes

Table 2: Sørensen's coefficient (SC), probability of detection (POD), false alarm ratio (FAR), and ignition time error values obtained for the SVM predictions.

Wildfire	SVM SC	SVM POD	SVM FAR	SVM ignition time error
Bobcat	0.77	0.95	0.35	-
Tennant	0.80	0.95	0.31	1 hour 56 minutes
Oak	0.77	> 0.99	0.38	25 minutes
Mineral	0.80	0.79	0.19	4 hour 13 minutes

accuracy of predicted fire perimeters. For predicted ignition times, in each case the cWGAN method is significantly more accurate than the SVM method.

5 Conclusion

Here a novel method for inferring the early-stage fire arrival time of a wildfire based on active fire satellite detections has been developed and tested on four California wildfires. A probabilistic approach is taken, enabling the ability to sample from the conditional distribution of fire arrival times given a measurement, which is accomplished using a conditional Wasserstein Generative Adversarial Network trained using solutions from WRF-SFIRE. The approach uniquely allows for quantification of prediction uncertainty and further allows physics to be inject into the predictions through the use of WRF-SFIRE solutions for training. Ultimately, this work enhances data assimilation capabilities for wildfire models, which are becoming increasingly important in our continually changing climate.

Acknowledgments and Disclosure of Funding

BS, DR, and AAO gratefully acknowledge support from ARO, USA grant W911NF2010050. KH was funded by NASA Disasters Program under grant 80NSSC19K1091. JM was funded by NASA grants 80NSSC19K1091 and 80NSSC22K1717. AK and AF were funded by NASA grants 80NSSC19K1091, 80NSSC22K1717, and 80NSSC22K1405, as well as NSF grants DEB-2039552 and IUCRC-2113931. JH is supported by NOAA under award NA22OAR4050672I. VC acknowledges support from the Center for Undergraduate Research at Viterbi Engineering (CURVE) program. The authors acknowledge the Center for Advanced Research Computing (CARC) at the University of Southern California, USA for providing computing resources that have contributed to the research results reported within this publication. The authors acknowledge the use of HPC infrastructure at the Center for Computational Mathematics, University of Colorado Denver, funded by NSF grant OAC-2019089.

References

[1] Anthony L Westerling, Hugo G Hidalgo, Daniel R Cayan, and Thomas W Swetnam. Warming and earlier spring increase western US forest wildfire activity. *Science*, 313(5789):940–943, 2006. doi: 10.1126/science.1128834.

[2] Philip E Dennison, Simon C Brewer, James D Arnold, and Max A Moritz. Large wildfire trends in the western United States, 1984–2011. *Geophysical Research Letters*, 41(8):2928–2933, 2014. doi: 10.1002/2014GL059576.

[3] A Park Williams, John T Abatzoglou, Alexander Gershunov, Janin Guzman-Morales, Daniel A Bishop, Jennifer K Balch, and Dennis P Lettenmaier. Observed impacts of anthropogenic climate change on wildfire in California. *Earth's Future*, 7(8):892–910, 2019. doi: 10.1029/2019EF001210.

[4] Yongqiang Liu, John Stanturf, and Scott Goodrick. Trends in global wildfire potential in a changing climate. *For. Ecol. Manage.*, 259(4):685–697, 2010. doi: 10.1016/j.foreco.2009.09.002.

[5] Dan Jaffe, William Hafner, Duli Chand, Anthony Westerling, and Dominick Spracklen. Interannual variations in PM2. 5 due to wildfires in the Western United States. *Environ. Sci. Technol.*, 42(8):2812–2818, 2008. doi: 10.1021/es702755v.

[6] Daoping Wang, Dabo Guan, Shupeng Zhu, Michael Mac Kinnon, Guannan Geng, Qiang Zhang, Heran Zheng, Tianyang Lei, Shuai Shao, Peng Gong, et al. Economic footprint of California wildfires in 2018. *Nat. Sustain.*, 4(3):252–260, 2021. doi: 10.1038/s41893-020-00646-7.

[7] Rosana Aguilera, Thomas Corringham, Alexander Gershunov, and Tarik Benmarhnia. Wildfire smoke impacts respiratory health more than fine particles from other sources: observational evidence from Southern California. *Nat. Commun.*, 12(1):1–8, 2021. doi: 10.1038/s41467-021-21708-0.

[8] Susan Solomon, Kimberlee Dube, Kane Stone, Pengfei Yu, Doug Kinnison, Owen B Toon, Susan E Strahan, Karen H Rosenlof, Robert Portmann, Sean Davis, et al. On the stratospheric chemistry of midlatitude wildfire smoke. *Proc. Natl. Acad. Sci.*, 119(10):e2117325119, 2022. doi: 10.1073/pnas.2117325119.

[9] Atoossa Bakhshaii and Edward A Johnson. A review of a new generation of wildfire–atmosphere modeling. *Can. J. For. Res.*, 49(6):565–574, 2019. doi: 10.1139/cjfr-2018-0138.

[10] Neil P Lareau and Craig B Clements. The mean and turbulent properties of a wildfire convective plume. *J. Appl. Meteor. Climatol.*, 56(8):2289–2299, 2017. doi: 10.1175/JAMC-D-16-0384.1.

[11] Neil P Lareau, NJ Nauslar, and John T Abatzoglou. The Carr Fire vortex: A case of pyrotornadogenesis? *Geophys. Res. Lett.*, 45(23):13–107, 2018. doi: 10.1029/2018GL080667.

[12] Wilfrid Schroeder and Louis Giglio. Visible Infrared Imaging Radiometer Suite (VIIRS) 375 m active fire detection and characterization algorithm theoretical basis document. *University of Maryland: Washington, DC, USA*, 2016.

[13] Jan Mandel, Jonathan D Beezley, Adam K Kochanski, Volodymyr Y Kondratenko, and Minjeong Kim. Assimilation of perimeter data and coupling with fuel moisture in a wildland fire–atmosphere dddas. *Procedia Computer Science*, 9:1100–1109, 2012. doi: 10.1016/j.procs.2012.04.119.

[14] J Mandel, Shai Amram, JD Beezley, Guy Kelman, AK Kochanski, VY Kondratenko, BH Lynn, B Regev, and Martin Vejmelka. Recent advances and applications of WRF–SFIRE. *Natural Hazards and Earth System Sciences*, 14(10):2829–2845, 2014. doi: 10.5194/nhess-14-2829-2014.

[15] Angel Farguell, Jan Mandel, James Haley, Derek V Mallia, Adam Kochanski, and Kyle Hilburn. Machine learning estimation of fire arrival time from Level-2 Active Fires satellite data. *Remote Sens.*, 13(11):2203, 2021. doi: 10.3390/rs13112203.

[16] Deep Ray, Javier Murgoitio-Esandi, Agnimitra Dasgupta, and Assad A Oberai. Solution of physics-based inverse problems using conditional generative adversarial networks with full gradient penalty. *arXiv preprint arXiv:2306.04895*, 2023. doi: 10.48550/arXiv.2306.04895.

[17] Jonas Adler and Ozan Öktem. Deep bayesian inversion. *arXiv preprint arXiv:1811.05910*, 2018. doi: 10.48550/arXiv.1811.05910.

[18] Deep Ray, Harisankar Ramaswamy, Dhruv V Patel, and Assad A Oberai. The efficacy and generalizability of conditional GANs for posterior inference in physics-based inverse problems. *arXiv preprint arXiv:2202.07773*, 2022. doi: 10.48550/arXiv.2202.07773.

[19] Paul H Greenfield, Woodrow Smith, and David C Chamberlain. Phoenix-the new forest service airborne infrared fire detection and mapping system. In *2nd Int. Wildland Fire Ecology and Fire Management Congress and the 5th Symposium on Fire and Forest Meteorology*, 2003.

A cWGAN architecture

Schematics of the cWGAN architecture used here are provided in Figure 3, with additional components included in Figure 4.

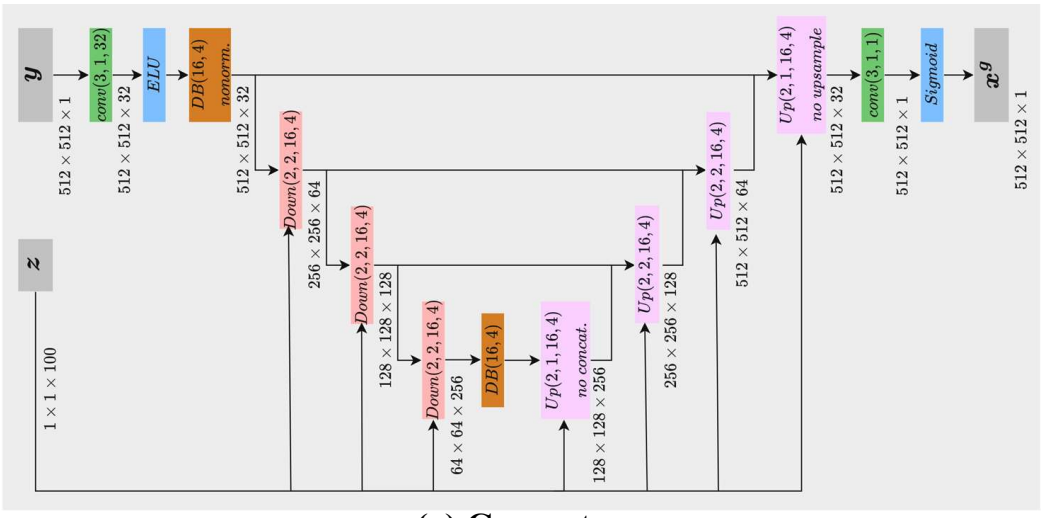

(a) Generator

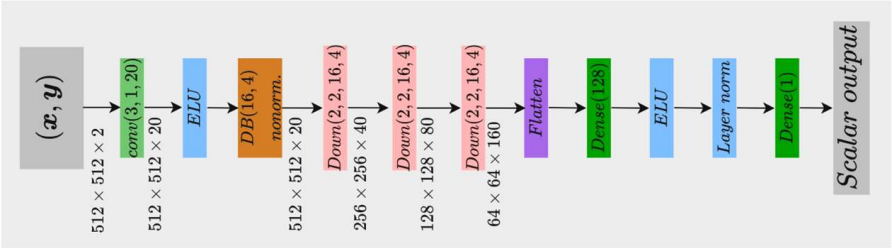

(b) Critic

Figure 3: Architecture of (a) generator and (b) critic used in cWGAN. The $Down$, Up and Dense (DB) blocks are described in Fig. 4. Latent vector $\boldsymbol{z}$ is injected into the generator using conditional instance normalization [18].

B Measurement operator

The measurement operator M used to generate an instance of measurement $\bar{\boldsymbol{\tau}}$ given a fire arrival time $\boldsymbol{\tau}$ was constructed to replicate the high resolution (375 m) VIIRS L2 AF data in the following steps:

1. Coarsen $\boldsymbol{\tau}^{(i)}$ obtained from data augmentation to a resolution of 375 m using nearest neighbor interpolation.
2. Select four measurement times ($t_i,\ i = 1, \cdots, 4$) from a uniform distribution between 2 hours and 48 hours and sort them in ascending order.
3. For each measurement time, t_i, create a time interval $(t_i - \delta^{(-)}, t_i)$ where $\delta^{(-)}$ is selected from $\mathcal{U}(6, 12)$, where $\mathcal{U}$ denotes the uniform probability distribution. If $t_i - \delta^{(-)} < 0$, set it to 0.
4. Create four copies of the coarsened $\boldsymbol{\tau}^{(i)}$, one for each time interval, and denote them by $\boldsymbol{\tau}_j^{(i)},\ j = 1, \cdots, 4$.
5. To each $\boldsymbol{\tau}_j^{(i)}$, apply a distinct knowledge mask that randomly eliminates 50% of the fire arrival time values. Set eliminated pixels to a background value.
6. For each $\boldsymbol{\tau}_j^{(i)}$, set fire arrival time pixels falling within the associated time interval $(t_i - \delta^{(-)}, t_i)$ to t_i. Set the remainder to a background value, to be assigned later.
7. Combine the four measurements into a single consolidated measurement by selecting $\boldsymbol{\tau}^{(i)} = \min_j(\boldsymbol{\tau}_j^{(i)})$ for each pixel.

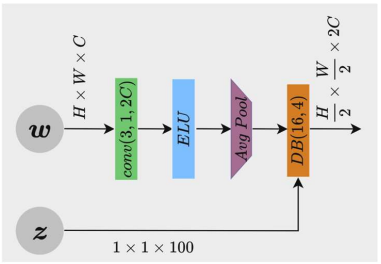

(a) Down(2,2,16,4)

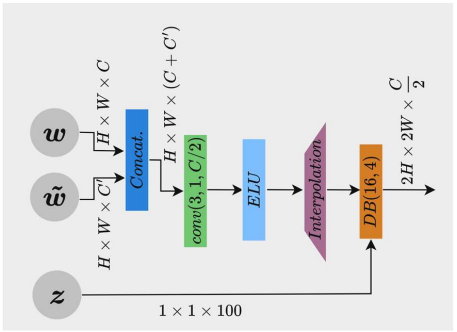

(b) Up(2,2,16,4)

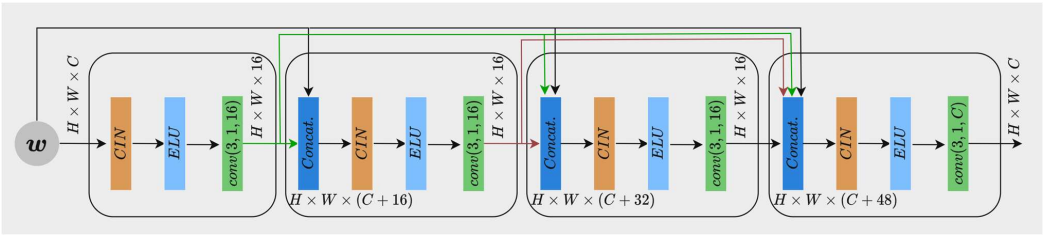

(c) DB(16,4)

Figure 4: Architecture of (a) down-sample block, (b) up-sample block, and (c) dense block, with the values $p = 2$, $q = 2$, $k = 16$, and $n = 4$ used for this work shown.

8. Eliminate three 3×3 km patches with locations selected at random. Set the values in these patches to the background value to emulate measurement obstruction.
9. Resample $\boldsymbol{\tau}^{(i)}$ back to the original size of 512×512 pixels.
10. Add $\delta \in \mathcal{U}(0, 24)$ hours to the arrival time and measurement pair $(\boldsymbol{\tau}^{(i)}, \bar{\boldsymbol{\tau}}^{(i)})$ to account for the fact that the ignition time is typically unknown.
11. Normalize the arrival time and measurement pair $(\boldsymbol{\tau}^{(i)}, \bar{\boldsymbol{\tau}}^{(i)})$ to be in the interval $[0, 1]$ by dividing it by 72 hours and setting the background value to 1.

C Training data samples

Sample data pairs from the training data set are shown in Fig. 5.

D VIIRS AF measurements and IR perimeters

To preprocess VIIRS 375 m L2 AF satellite data to be used as input to the cWGAN, a domain of interest with size 30.72×30.72 km which is approximately centered on a desired wildfire is selected. The domain is discretized based on latitude and longitude coordinates and cells corresponding to AF detection locations are assigned a value based on the measurement time. Values assigned to activated cells are the number of hours since the start of the day on which ignition occurred. The measurements are then normalized using a value of 72 hours, following which remaining cells are assigned a background value of 1, putting the measurements in the range $[0, 1]$, following the format of the training data. This is done for all measurements available within the first 48 hours of a fire, being sure to assign the earliest available measurement time for cells that correspond to detections in more than one satellite measurement.

Shown in Fig. 6 are two measurements $\bar{\boldsymbol{\tau}}$ per fire, corresponding to the two confidence intervals. For each fire, one IR fire extent measurement made within the first 48 hours after ignition is overlaid on the measurement images for comparison. The times of the perimeter measurements are indicated in the plot labels of Fig. 6 in hours and minutes (HH:MM format) from the start of ignition day.

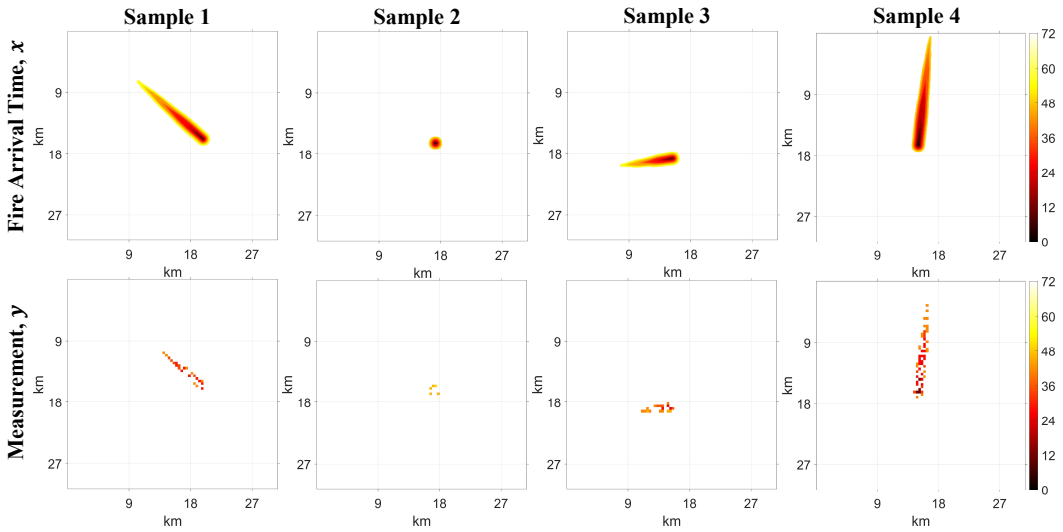

Figure 5: Sample data pairs from the training set, with true fire arrival times τ in the first row and corresponding measurements $\bar{\tau}$ in the second row. Here fire arrival time values represent hours from the start of the day on which ignition occurred.

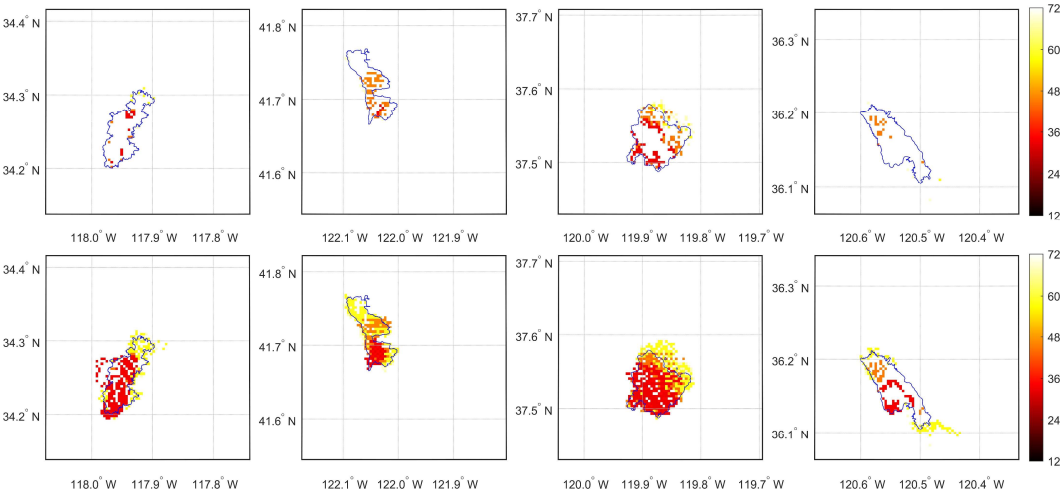

Figure 6: Measurements $\bar{\tau}$ after preprocessing of VIIRS 375 m L2 AF data for the Bobcat, Tennant, Oak, and Mineral fires, in left to right order. The first row contains high confidence detections only and the second row contains high and nominal confidence detections. AF detection colors indicate the measurement time, taken as the number of hours after the start of the ignition day. IR fire extent perimeters are additionally included, with measurement times listed in plot labels in HH:MM format, again as the number of hours after the start of ignition day. All measurements are geolocated, with longitude and latitude indicated.