remote sensing

MDPI

Article

Generative Algorithms for Wildfire Progression Reconstruction from Multi-Modal Satellite Active Fire Measurements and Terrain Height

Bryan Shaddy [1,*], Brianna Binder [1], Agnimitra Dasgupta [1], Haitong Qin [2], James Haley [3], Angel Farguell [4], Kyle Hilburn [3], Derek V. Mallia [5], Adam Kochanski [4], Jan Mandel [6] and Assad A. Oberai [1]

[1] Department of Aerospace and Mechanical Engineering, University of Southern California, Los Angeles, CA 90089, USA; bjbinder@usc.edu (B.B.); adasgupt@usc.edu (A.D.); aoberai@usc.edu (A.A.O.)
[2] Department of Mathematics, University of Southern California, Los Angeles, CA 90089, USA; qinmicha@usc.edu
[3] Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins, CO 80521, USA; james.haley@colostate.edu (J.H.); kyle.hilburn@colostate.edu (K.H.)
[4] Wildfire Interdisciplinary Research Center, San Jose State University, San Jose, CA 95192, USA; angel.farguellcaus@sjsu.edu (A.F.); adam.kochanski@sjsu.edu (A.K.)
[5] Department of Atmospheric Sciences, University of Utah, Salt Lake City, UT 84112, USA; derek.mallia@utah.edu
[6] Department of Mathematical and Statistical Sciences, University of Colorado Denver, Denver, CO 80204, USA; jan.mandel@ucdenver.edu
* Correspondence: bshaddy@usc.edu

Highlights

What are the main findings?

- Conditional generative algorithms trained on simulations of historic wildfires may be used to effectively reconstruct the early-time progression of wildfires given satellite active fire measurements and terrain height data.
- When applied to real wildfires, generated fire progression estimates compare favorably to ground-truth high resolution infrared perimeters measured via aircraft, with the ability to gather additional information about model uncertainty from generated samples.

What are the implications of the main findings?

- Once obtained, fire progression estimates may be used to perform data assimilation, wherein the estimated fire state is used to initialize subsequent wildfire spread forecasts. Fire progression estimates are additionally useful for providing situational awareness to wildfire stakeholders.
- The work developed here additionally demonstrates a framework for how critical characteristics affecting wildfire spread, such as terrain, may be used to improve estimates of a wildfire's state.

Abstract

Wildfire spread prediction models, including even the most sophisticated coupled atmosphere–wildfire models, diverge from observed wildfire progression during multi-day simulations, motivating the need for measurement-based assessments of wildfire state and improved data assimilation techniques. Data assimilation in the context of coupled atmosphere–wildfire models entails estimating wildfire progression history from observations and using this to obtain initial conditions for subsequent simulations through a spin-up process. In this study, an approach is developed for estimating fire progression history from VIIRS active fire measurements, GOES-derived ignition times, and terrain height data. The approach utilizes a conditional Wasserstein Generative Adversarial Network trained on simulations of historic wildfires from the coupled atmosphere–wildfire

Check for updates

Academic Editor: Nikos Koutsias

Received: 23 October 2025
Revised: 24 December 2025
Accepted: 8 January 2026
Published: 10 January 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

model WRF-SFIRE, with corresponding measurements for training obtained through the application of an approximate observation operator. Once trained, the cWGAN leverages measurements of real fires and corresponding terrain data to probabilistically generate fire progression estimates that are consistent with the WRF-SFIRE solutions used for training. The approach is validated on five Pacific US wildfires, and results are compared against high-resolution perimeters measured via aircraft, finding an average Sørensen–Dice coefficient of 0.81. The influence of terrain data on fire progression estimates is also assessed, finding an increased contribution when measurements are uninformative.

Keywords: conditional generative models; deep learning; fire monitoring; active fires satellite data; data assimilation; model initialization; Bayesian methods

1. Introduction

1.1. State Estimation for Wildfires

Climate trends and variability, including drier summers and increased heat, have contributed to higher wildfire occurrence rates and greater fire growth potential during the late 20th and early 21st centuries [1,2]. Predictions for the remainder of this century indicate even larger wildfires with increased intensity, accompanied by a lengthening of wildfire season duration for many regions globally [3,4]. Negative impacts associated with increased wildfire presence are compounded by growing numbers of residential homes on the wildland–urban interface in the United States, with current estimates putting the number at around 50 million and increasing by approximately 350,000 houses per year, leading to greater potential for communities to be directly affected by wildfires [5]. Due in part to the increasing proximity of wildfires to populated areas, the economic impact of wildfires has also grown significantly. Damage from the 2018 California wildfires was estimated at nearly USD 150 billion , including capital losses, health costs, and indirect losses [6]. More recently, initial estimates place the economic impact of the Los Angeles wildfires occurring during the first two weeks of 2025 at around USD 250 billion [7].

Wildfire spread is influenced by the interaction of fire with fuels, topography, and weather conditions. With the increasing size and intensity of modern wildfires, their interactions with the atmosphere are growing in complexity and becoming an even more critical aspect of predicting fire behavior [8–12]. To capture interactions between wildfires and the atmosphere, in addition to topography and fuel, state-of-the-art models such as WRF-SFIRE consider two-way atmosphere–wildfire coupling. In these coupled models, heat released to the atmosphere by a wildfire is computed and used to modify local meteorological conditions, which are then fed back into the spread of the fire [13,14]. However, even forecasts from these complex models aiming to capture all critical aspects of wildfire spread ultimately diverge from the true state of a wildfire over the course of a one- to two-day simulation. This in turn underscores the need for data assimilation, which integrates measurements of ongoing wildfires to enhance prediction accuracy [15].

Beyond its use for initialization of coupled atmosphere–wildfire models, estimates of the past progression and current state of a wildfire also prove useful for several additional wildfire management tasks. They provide situational awareness to fire managers and other stakeholders addressing ongoing fires, enabling the contextualization of wildfire activity through mapping against terrain and other informative variables [16,17]. From fire spread history estimates, information about the spread rate, intensity, and potential growth may also be gathered, aiding identification of highly active regions of a fire and assisting the development of appropriate mitigation strategies [18,19]. The graphical nature of fire state

estimates may also be leveraged within early-warning systems, allowing for more effective warnings and safer, more efficient evacuation of citizens [20].

Common systems for monitoring wildfire spread include satellite-based sensors, which are deployed on polar orbiting and geostationary satellites. These satellite systems utilize thermal imaging to provide active fire (AF) detections, which indicate that the time and location an active fire was identified for a given area. The Visible Infrared Imaging Radiometer Suite (VIIRS), which is carried on polar-orbiting satellites including Suomi NPP, NOAA-20, and NOAA-21, provides AF detection with a spatial resolution of 375 m roughly 2–4 times per day at mid-latitudes [21]. The geostationary GOES satellite series, serving operationally as GOES-East and GOES-West, provides similar data through the Fire/Hot Spot Characterization product, producing measurements every 5 min, though at a coarser spatial resolution of 2–4 km [22]. AF detections from these satellite systems are useful for providing information about wildfire progression; however, the data are at spatial and temporal scales that are orders of magnitude coarser than those of modern wildfire spread models, complicating their use for performing data assimilation. Data from these systems also tends to be sparse and noisy due to artifacts caused by sun glint, clouds, topography, and in some cases, hot smoke or embers. However, even with these limitations, AF data serves as a useful source of measurements for performing data assimilation, where it may be used to infer a smooth and finely sampled version of the fire state.

Two approaches have proven useful for performing data assimilation in coupled atmosphere–wildfire models. The first involves only using a series of measurements of a wildfire during its initial spread to estimate the fire and atmosphere states; this can be thought of as initial state estimation [23,24]. The second approach involves using both current predictions of the wildfire and atmosphere states from a forward prediction model and a newly obtained set of measurements to determine updated state estimates; this can be thought of as sequential state estimation [25]. Both approaches result in a set of conditions which may be used to perform subsequent forecasts.

1.2. The Role of Fire Arrival Time

For initial and sequential state estimation, it is critical to recognize that during data assimilation, both wildfire and atmosphere variables must be updated in a self-consistent manner. To accomplish this, Mandel et al. [24] recognized that if the historic spread of a wildfire is determined during the assimilation process, it may be used to prescribe fire progression within a coupled atmosphere–wildfire model, thereby generating the fire heat and moisture fluxes needed to produce the atmospheric state consistent with the fire. Kochanski et al. [26] later demonstrated the importance of this type of gradual ignition for developing physically-consistent atmospheric states, which in turn enables accurate prediction of fire progression and pyroconvection in subsequent forecasts. This process of gradual ignition is referred to as spin-up, and simplifies the state estimation problem to one concerned solely with estimating fire progression.

The fire arrival time, which is the time from ignition, the fire front reaches a given location, is a convenient means to represent wildfire spread. It provides a succinct representation of fire progression in a single field which can be represented on a spatial grid [24,27]. Therefore, for either data assimilation problem, we are interested in estimating fire arrival times for some time interval, either from ignition to the current time or the previous to the current data assimilation steps. Once determined, the fire arrival time may be used to prescribe wildfire progression within a coupled model to develop the corresponding atmospheric state. Subsequent forecasts may then be initialized from these fire and atmospheric states.

1.3. Prior Work and Proposed Developments

Over the last decade, several methods have been developed to estimate the state of actively burning wildfires [27–35]. These approaches typically fall into two categories: (a) approaches which provide the perimeter of an active fire, often on daily or twice-daily timescales, and (b) approaches providing information about the fine-scale progression of a fire, typically in the form of the fire arrival time. A key distinction between these approaches is their ability to provide sufficient information on wildfire progression to spin-up the atmospheric state of coupled atmosphere–wildfire models, a process shown to be critical to the accuracy of simulations from these models [26]. Beyond the temporal resolution of these approaches, spatial resolution can also play a role in their ability to provide sufficiently detailed fire progression estimates for accurate initialization of coupled models.

A number of geospatial interpolation schemes for estimating fire arrival times from satellite measurements have been developed, though in the context of coupled model initialization limitations of these approaches arise due to their spatial and temporal resolution [30–32]. The Fire Events Data Suite (FEDS) is another system which has become quite popular in recent years for mapping wildfire progression [33]. The FEDS system uses an object-based system for tracking wildfire progression based on VIIRS data, providing wildfire perimeters twice-daily. However, because FEDS only provides fire progression information in the form of twice-daily perimeters, further interpolation must be performed to obtain fire progression information sufficient for initializing coupled models. Following the work of Kochanski et al. [26] on initialization methods for coupled atmosphere–wildfire models using aircraft-measured fire perimeters and overcoming the coarse temporal frequency of many fire state estimates, Clough et al. [34] explored the use of consecutive fire perimeters to estimate fire arrival times using linear interpolation between perimeters.

Recent years have additionally seen a number of works using machine learning approaches for fire arrival time estimation. In Farguell et al. [28], the authors use a support vector machine (SVM) to estimate fire arrival times using satellite AF detections and clear ground pixels, providing a robust method for reconstructing fire progression history, though at the expense of slow processing times. In McCarthy et al. [35], the use of a convolutional neural network (CNN) to downscale GOES data under consideration of terrain, vegetation, and land use information to obtain wildfire progression estimates with a spatial resolution of 375 m and at a temporal resolution of 5 min was explored. However, this model relies on GOES data to obtain a high temporal resolution, but at the cost of spatially coarse input data, which makes the estimation of the spatial progression of a fire challenging. Shaddy et al. [29] later proposed the use of a conditional Wasserstein Generative Adversarial Network (cWGAN) trained using WRF-SFIRE solutions to probabilistically reconstruct fire arrival times from AF measurements. This approach allowed for uncertainty quantification during inference and the inclusion of physics through training data. However, it used only data from the VIIRS system and did not include the effect of additional fields like terrain or fuel which are known to influence wildfire spread.

The present work extends the cWGAN-based approach for initial state estimation of Shaddy et al. [29] in several significant ways. First, in addition to VIIRS data, it takes advantage of AF data from GOES to estimate the ignition time for use during fire arrival time inference. Second, it demonstrates how the inferred arrival time can account for variables like terrain, fuel and weather by conditioning the inference on these variables. As an initial step in this direction, in the current work, we condition only on the terrain map. Finally, for training the cWGAN, we use considerably more realistic data. In particular, we use WRF-SFIRE simulation results of historical wildfires occurring in CONUS during 2023, thus allowing for more complex and realistic wildfire spread physics to be included in the fire arrival time estimates. This is in contrast to [29], where the training data comprised

arrival times for simulated wildfires occurring on a flat terrain with a uniform fuel type and simplified atmospheric conditions.

Our approach begins with the generation of wildfire spread realizations using WRF-SFIRE, where we simulate 48 h of spread for named wildfires occurring in CONUS during 2023. These wildfire spread realizations provide target fire arrival times for training, along with their corresponding terrain height maps. To generate measurements for training, we then apply an approximate observation operator (detailed in Appendix A) to the fire arrival times to simulate measurements from the VIIRS system with ignition times derived from GOES. The 3-tuple comprising arrival times, VIIRS measurements, and terrain are treated as random variables, and are used to train a cWGAN. Once trained, the cWGAN is used to generate samples of the fire arrival time conditioned on VIIRS AF measurements and terrain (see Figure 1 for a schematic description). The samples thus obtained are used to compute statistics of interest, including the mean and standard deviation of arrival time maps, providing a best guess estimate and a measure of uncertainty. The performance of the resulting algorithm is quantified by applying it to historical wildfires not included in the training data.

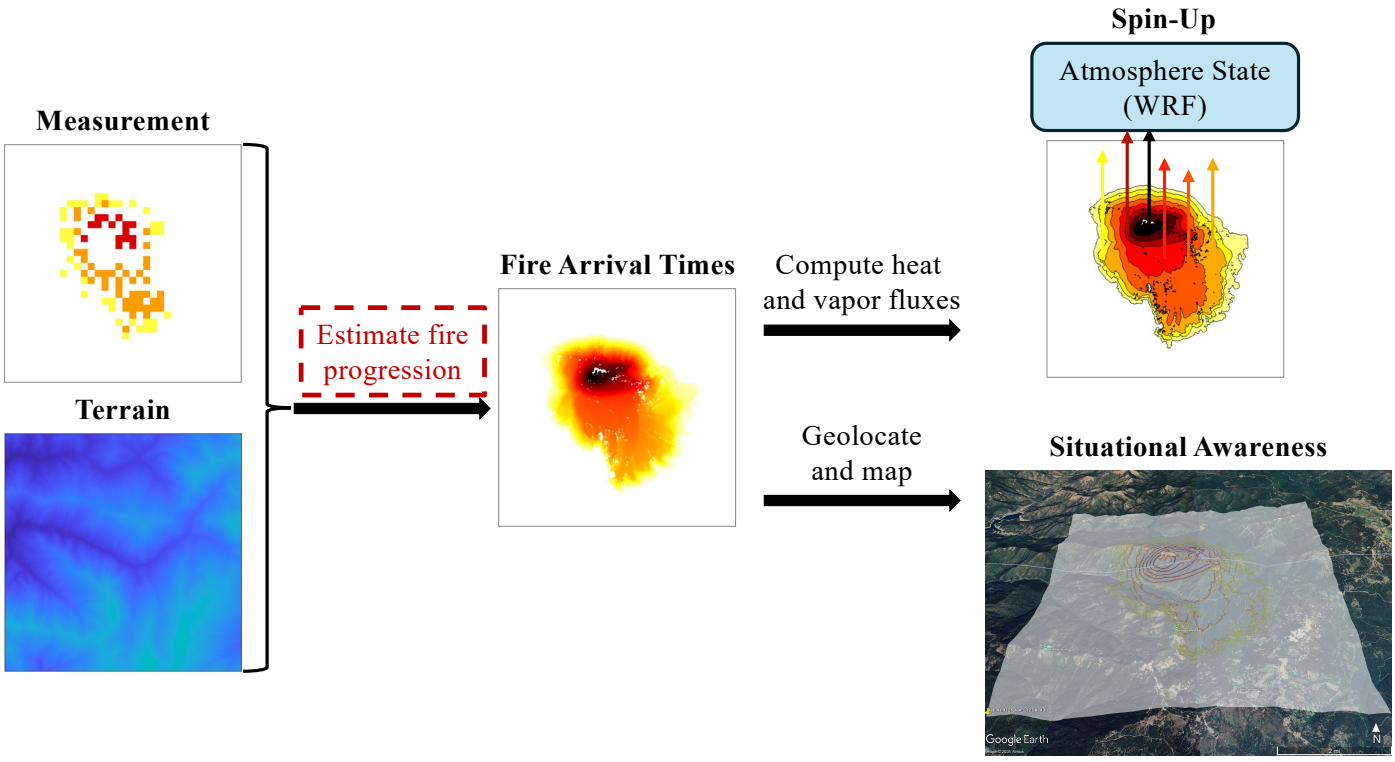


Figure 1. Overview of the fire arrival time estimation task addressed here, along with downstream applications including spin-up of the atmospheric state within coupled atmosphere–wildfire models (data assimilation) and mapping of fire progression for situational awareness.

For this work, we chose to utilize cWGANs for their straight-forward and timely sampling approach, which requires just one forward pass through the generator to produce a single sample. While other more recent generative modeling algorithms, such as diffusion models, have also been shown to produce accurate results in a number of tasks, the generation of samples requires many forward passes, in addition to the integration of a differential equation [36]. The cost of this computation additionally grows as the problem dimensionality increases, becoming a drawback for tasks which are both time-sensitive and high-dimensional, such as wildfire state estimation. Work is being conducted to improve the sampling efficiency of diffusion models, which will likely make their use more appealing for wildfire state estimation in the near future [37–40].

The main contributions and innovations of this work are summarized in the following list of key advancements:

1. A novel terrain-conditioned fire arrival time inference framework is developed.
2. Multi-modal satellite observations from VIIRS and GOES, combined with terrain data, are used as the basis for fire arrival time estimates.
3. A training paradigm grounded in WRF-SFIRE simulations of historical wildfires is utilized to incorporate realistic physics into arrival time estimates.
4. Probabilistic estimates are used to quantify uncertainty.

The remainder of the manuscript is structured as follows. In Section 2.1, we describe and formulate the problem of probabilistically inferring fire arrival times from satellite measurements and terrain height data. Thereafter, in Section 2.2, we discuss how wildfire spread solutions obtained using WRF-SFIRE are used to generate the training data. Following this, in Section 2.3, we describe the cWGAN model and the training process. In Section 3, we present results for five Pacific US wildfires and validate the predicted fire perimeters against high-resolution perimeters measured using aircraft-mounted infrared sensors. We also examine the effect of the terrain, which is the conditioning variable, on the arrival time reconstruction. Section 4 then provides a discussion of the results. Lastly, in Section 5, we provide concluding remarks including a discussion of limitations and possible future developments.

2. Materials and Methods

2.1. Problem Formulation

Let the vector of fire arrival times be denoted by the random vector $\boldsymbol{T}$ with the probability density $P_{\boldsymbol{T}}$. Samples $\boldsymbol{\tau}$ of $\boldsymbol{T}$ consist of N components τ_i and therefore $\boldsymbol{\tau} \in \Omega_{\boldsymbol{\tau}} \equiv \mathbb{R}^N$. The components of $\boldsymbol{\tau}$ are measured on a 12.8 km $\times$ 12.8 km domain with a resolution of 25 m, and there are $N = 512 \times 512$ such components. An observation operator M defined by the mapping $M : \Omega_{\boldsymbol{\tau}} \to \Omega_{\bar{\boldsymbol{\tau}}}$ may be applied to fire arrival times $\boldsymbol{\tau}$ to generate coarse, sparse and noisy measurements of arrival time $M(\boldsymbol{\tau}) = \bar{\boldsymbol{\tau}}$, with their own probability density $P_{\bar{\boldsymbol{T}}}$. The measured arrival times are defined on the same grid, and therefore $\bar{\boldsymbol{\tau}} \in \Omega_{\bar{\boldsymbol{\tau}}} \equiv \mathbb{R}^N$. Here, the mapping M is defined through a set of steps (described in Appendix A) that approximate the measurements generated by satellite active fire products. The terrain height vector is described on the same grid and is represented by the random vector $\boldsymbol{H}$ with a probability density $P_{\boldsymbol{H}}$ and samples $\boldsymbol{h} \in \Omega_{\boldsymbol{h}} \equiv \mathbb{R}^N$.

Here, we are interested in solving for the inverse mapping which transforms measurements $\bar{\boldsymbol{\tau}}$ to fire arrival times $\boldsymbol{\tau}$. Further, we would like to condition this inference on terrain height $\boldsymbol{h}$. We note that, unlike the forward mapping, the inverse mapping is ill-posed because a single measurement–terrain pair $(\bar{\boldsymbol{\tau}}, \boldsymbol{h})$ can correspond to many different fire arrival times $\boldsymbol{\tau}$. To address the ill-posed nature of the inverse map we propose a probabilistic approach wherein we quantify the conditional distribution $P_{\boldsymbol{T}|\bar{\boldsymbol{T}},\boldsymbol{H}}$ such that for a given measurement–terrain pair $(\bar{\boldsymbol{\tau}}, \boldsymbol{h})$ we are able to draw samples of $\boldsymbol{\tau}$ from this conditional distribution. To accomplish this, we use a cWGAN trained on samples of $(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}, \boldsymbol{h})$ drawn from the joint distribution $P_{\boldsymbol{T},\bar{\boldsymbol{T}},\boldsymbol{H}}$. Once trained, for a given $(\bar{\boldsymbol{\tau}}, \boldsymbol{h})$ the generator of the cWGAN produces samples of $\boldsymbol{\tau}$ from the learned conditional distribution $P^{\boldsymbol{g}}_{\boldsymbol{T}|\bar{\boldsymbol{T}},\boldsymbol{H}} \approx P_{\boldsymbol{T}|\bar{\boldsymbol{T}},\boldsymbol{H}}$. These samples are used to compute statistics such as the mean and standard deviation of the arrival time.

To generate training data, first samples of terrain height $\boldsymbol{h}$ are drawn from the marginal distribution $P_{\boldsymbol{H}}$. Following this, WRF-SFIRE is employed to draw corresponding samples of fire arrival times $\boldsymbol{\tau}$ from the conditional distribution $P_{\boldsymbol{T}|\boldsymbol{H}}$ by running wildfire progression simulations for a given $\boldsymbol{h}$. The measurement operator M is then applied to samples of $\boldsymbol{\tau}$ to provide corresponding samples $\bar{\boldsymbol{\tau}}$. Thus, tuples of $(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}, \boldsymbol{h})$ from the joint distribution

$P_{\boldsymbol{T},\bar{\boldsymbol{T}},\boldsymbol{H}}$ are obtained. With samples from the joint distribution obtained, the cWGAN may be trained to sample from the conditional distribution $P_{\boldsymbol{T}|\bar{\boldsymbol{T}},\boldsymbol{H}}$. The following section includes a detailed discussion of the training data generation process.

2.2. Training Data

To generate training data consisting of tuples of fire arrival times $\boldsymbol{\tau}$, measurements $\bar{\boldsymbol{\tau}}$, and terrain height $\boldsymbol{h}$, we employ the fully coupled atmosphere–wildfire model WRF-SFIRE, which couples the Weather Research and Forecasting model with a fire spread parameterization (SFIRE). SFIRE is based on the level-set method and uses the semi-empirical Rothermel rate of spread model to propagate fire growth [14,41]. WRF-SFIRE simulations were run utilizing the wrfxpy system for automated simulation initiation, which configures the simulation domain, collects and processes data for initial conditions and boundary conditions, and launches the simulation [42]. Corresponding measurements are derived from WRF-SFIRE solutions through application of an approximate observation operator. The following sections outline the use of the WRF-SFIRE model to generate wildfire spread solutions and describe how training data was subsequently constructed from these solutions.

2.2.1. WRF-SFIRE Simulations

A total of 140 WRF-SFIRE simulations of wildfires that occurred in the CONUS during 2023 were used to train our model, with 12 additional WRF-SFIRE simulations used for validation and hyperparameter tuning. Wildfire incidents were selected manually based on simulated extents after 24 h, excluding those with potential to grow larger than our computational domain or those deemed too small to be representative. To initialize the WRF-SFIRE simulations, ignition locations and times were derived from GOES and VIIRS active fire detections, delivered by NOAA's Next Generation Fire System (NGFS). The algorithm used by the NGFS clusters detections into incidents which are then associated with reported fires cataloged by the National Interagency Fire Center (NIFC). The ignition time is estimated to be the time of the first detection belonging to a cluster of detections associated with a fire event. The ignition location is estimated from the average location of all VIIRS detections found within the footprint of the initial GOES detection. If no VIIRS detections are observed within the first four hours after the initial GOES detection, an average of the GOES detection locations is used.

Simulations were run using a domain of size 30 km $\times$ 30 km, with an atmospheric grid of size 30 $\times$ 30 with 1 km resolution. The atmospheric grid additionally contained 40 vertical levels, with a pressure top of roughly 5000 Pa (approximately 20 km). The wildfire state was tracked on a grid of size 1200 $\times$ 1200 with 25 m resolution, using a grid refinement ratio between the atmospheric grid and fire grid of 40:1. Domains were constructed with ignition locations approximately in their center. Initial and boundary conditions were obtained from the North American Mesoscale (NAM) Forecast System Grid 227, which was collected at 3 h intervals and provided data with 5 km resolution. NAM provides data for a variety of meteorological variables such as wind, temperature, pressure, and humidity, and was also used to provide boundary conditions for the fuel moisture model. Fuel category data came from the 2019 LandFire 13 Anderson Fire Behavior Fuel Models (FBFM13) dataset and terrain height data was obtained from the USGS National Elevation Dataset (NED). Following simulation initialization, the model projected fire spread over a 48 h period, resulting in the creation of the fire arrival time maps $\boldsymbol{\tau}$ used for training data generation.

We note that while WRF-SFIRE itself may deviate from the true progression of a wildfire over the course of a 24–48 h simulation due to the chaotic nature of wildfires, as would

any wildfire spread prediction model, we do not rely on the results from these models alone. Rather, we ground our predictions in truth by also conditioning on satellite measurements. Thus, the WRF-SFIRE simulations are used to learn how to interpolate between satellite detections, meaning that even if these simulations are not perfect, reasonable results are still obtainable.

2.2.2. Construction of Training Data

Once fire arrival time maps $\boldsymbol{\tau}$ and corresponding terrain height maps $\boldsymbol{h}$ are obtained, arrival time values are adjusted to ensure ignition occurs at time zero, after which data augmentation is performed to the tuples $(\boldsymbol{\tau}, \boldsymbol{h})$. Data augmentation was performed by first rotating randomly between 0 and 360 degrees, then translating randomly in a box of size 1 km $\times$ 1 km, and finally cropping to a size of 512 $\times$ 512 pixels (12.8 km $\times$ 12.8 km). These data augmentation steps were performed 25 times per each WRF-SFIRE fire arrival time solution, following which the observation operator M was applied to each augmented solution 5 times, creating 5 unique measurements per augmented target fire arrival time and terrain pair. In total this leads to 17,500 tuples $(\boldsymbol{\tau}^{(i)}, \bar{\boldsymbol{\tau}}^{(i)}, \boldsymbol{h}^{(i)})$ for training and an additional 1500 tuples for validation. Figure 2 shows sample tuples from the training data set. Prior to training, fire arrival times $\boldsymbol{\tau}^{(i)}$ and measurements $\bar{\boldsymbol{\tau}}^{(i)}$ are normalized by dividing by 48 h, and terrain height $\boldsymbol{h}$ is normalized by subtracting the minimum value and dividing by 3000 m, ensuring that all values are in the range $[0, 1]$.

The observation operator used here to construct training data was designed to sample artificial measurements $\bar{\boldsymbol{\tau}}^{(i)}$ from a probability distribution approximating the distribution of measurements that may be obtained from 375 m Level-2 VIIRS AF data resulting from multiple satellite overpasses, with ignition times gathered from GOES Fire/Hot Spot Characterization data, for a given fire arrival time map $\boldsymbol{\tau}^{(i)}$. Details of the measurement operator approximation are provided in Appendix A. By simulating critical characteristics of active fire data, resulting artificial measurements contain the same sparse, coarse, and noisy patterns as observed for true measurements, though alternative observation operators accounting for things such as fire heat flux, which are not considered here, may be devised (e.g., [43]).

We additionally note that data augmentation was performed so as to reach a number of training samples on the order of 10–20 thousand, often shown to be sufficient in image-based generative modeling tasks, while the total number of WRF-SFIRE solutions used to produce training data may be a limitation which could lead to overfitting, the model state was evaluated on a separate validation dataset throughout the training process, and the model state resulting in the best overall performance was selected, ensuring that overfitting was avoided.

2.3. Conditional Wasserstein Generative Adversarial Network (cWGAN)

To solve the target inverse problem of inferring fire arrival times $\boldsymbol{\tau}$ from measurements $\bar{\boldsymbol{\tau}}$ and terrain height $\boldsymbol{h}$, we employ a cWGAN. The cWGAN comprises two subnetworks, a generator $\boldsymbol{g}$ and a critic d, which are given by the mappings $\boldsymbol{g} : \Omega_{\bar{\tau}} \times \Omega_h \times \Omega_z \to \Omega_\tau$ and $d : \Omega_\tau \times \Omega_{\bar{\tau}} \times \Omega_h \to \mathbb{R}$, respectively. Here, $\boldsymbol{z} \in \Omega_z \subset \mathbb{R}^{N_z}$ is a latent variable modeled using the random variable $\boldsymbol{Z}$ sampled from a multivariate Gaussian distribution $P_{\boldsymbol{Z}}$. Once trained, given a measurement–terrain pair $(\bar{\boldsymbol{\tau}}, \boldsymbol{h})$, the generator $\boldsymbol{g}$ produces samples $\boldsymbol{\tau}^g = \boldsymbol{g}(\bar{\boldsymbol{\tau}}, \boldsymbol{h}, \boldsymbol{z}), \boldsymbol{z} \sim P_{\boldsymbol{Z}}$ from the learned conditional distribution $P^{\boldsymbol{g}}_{\boldsymbol{T}|\bar{\boldsymbol{T}},\boldsymbol{H}}(\boldsymbol{\tau}|\bar{\boldsymbol{\tau}}, \boldsymbol{h})$.

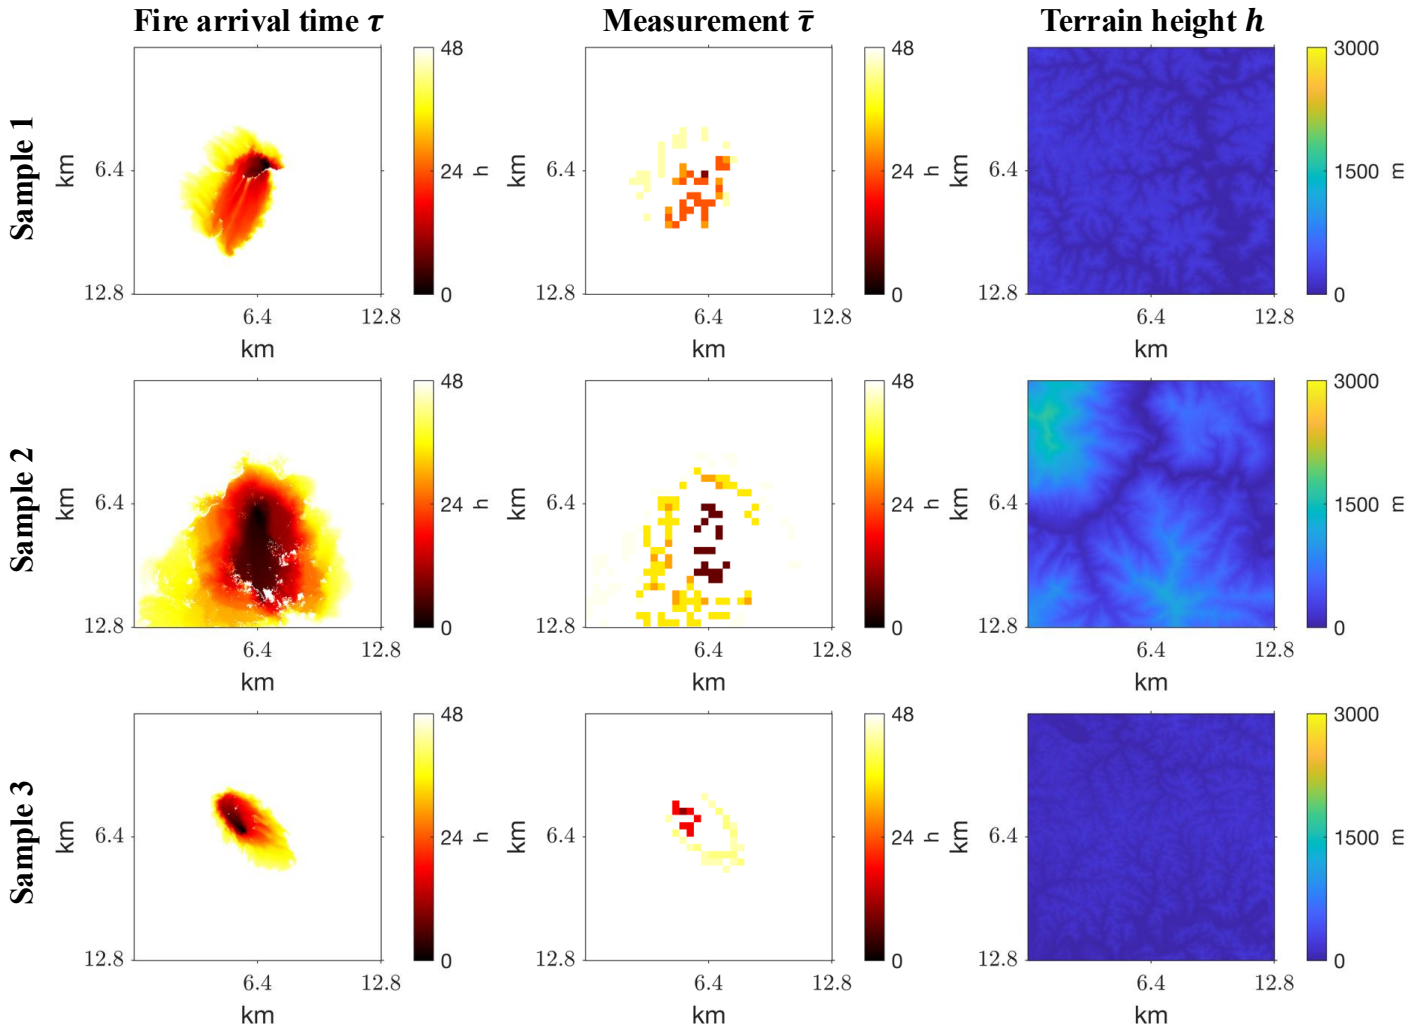


Figure 2. Training data samples, with target fire arrival time solutions $\boldsymbol{\tau}$ in the first column, measurements $\bar{\boldsymbol{\tau}}$ in the second column, and terrain height $\boldsymbol{h}$ in the third column.

Training of the cWGAN requires that the Wasserstein-1 distance between the learned distribution $P^{g}_{T|\bar{T},H}(\boldsymbol{\tau}|\bar{\boldsymbol{\tau}},\boldsymbol{h})$ and the true conditional distribution $P_{T|\bar{T},H}(\boldsymbol{\tau}|\bar{\boldsymbol{\tau}},\boldsymbol{h})$ be minimized. To accomplish this, a min-max problem is solved using the objective function

$$\mathcal{L}(d,g) = \mathop{\mathbb{E}}_{(\boldsymbol{\tau},\bar{\boldsymbol{\tau}},\boldsymbol{h})\sim P_{T,\bar{T},H}} [d(\boldsymbol{\tau},\bar{\boldsymbol{\tau}},\boldsymbol{h})] - \mathop{\mathbb{E}}_{\substack{\boldsymbol{\tau}^{g}\sim P^{g}_{T|\bar{T},H} \\ (\bar{\boldsymbol{\tau}},\boldsymbol{h})\sim P_{\bar{T},H}}} [d(\boldsymbol{\tau}^{g},\bar{\boldsymbol{\tau}},\boldsymbol{h})], \tag{1}$$

where the optimal generator and critic, denoted by g^* and d^*, respectively, are found as

$$(d^*,g^*) = \arg\min_{g} \arg\max_{d} \mathcal{L}(d,g). \tag{2}$$

The optimal g^* found by solving the min-max problem given by Equation 2 can then be used to approximate the true conditional distribution, assuming the critic is 1-Lipschitz in all of its arguments [44]. Selecting a generator with a sufficiently large number of learnable parameters, for any continuous, bounded function $\ell(\boldsymbol{\tau})$ defined on Ω_{τ} and for $\epsilon > 0$, it can be shown that

$$\left| \mathop{\mathbb{E}}_{\boldsymbol{\tau}\sim P_{T|\bar{T},H}} [\ell(\boldsymbol{\tau})] - \mathop{\mathbb{E}}_{\boldsymbol{\tau}^{g}\sim P^{g^*}_{T|\bar{T},H}} [\ell(\boldsymbol{\tau}^{g})] \right| < \epsilon. \tag{3}$$

Therefore, the trained generator may be used to approximate statistics from the true conditional distribution as

$$\mathop{\mathbb{E}}_{\boldsymbol{\tau}\sim P_{T|\bar{T},\boldsymbol{H}}}[\ell(\boldsymbol{\tau})] \approx \frac{1}{K}\sum_{i=1}^{K}\ell(\boldsymbol{g}^*(\bar{\boldsymbol{\tau}},\boldsymbol{h},\boldsymbol{z}^{(i)})), \quad \boldsymbol{z}^{(i)} \sim P_{\boldsymbol{Z}}, \tag{4}$$

where by setting $\ell(\boldsymbol{\tau}) = \boldsymbol{\tau}$ the pixel-wise mean prediction may be computed, and by setting $\ell(\boldsymbol{\tau}) = (\boldsymbol{\tau} - \mathbb{E}[\boldsymbol{\tau}])^2$ the pixel-wise variance may be computed. For additional background we refer readers to [44–46].

Details of the cWGAN network architecture utilized here are provided in Appendix B. We note that here we have chosen to utilize a U-Net-based architecture for our generator, as opposed to a transformer-based architecture, which has become a popular alternative, primarily due to the limited training data available. The amount of training data available for this work is likely insufficient to train transformer-based models due to the large number of parameters they typically contain.

Training cWGAN

The cWGAN was implemented in PyTorch version 2.3.0 and trained using 2 A100 GPUs for 160 epochs. A batch size of 16 was used, along with the Adam optimizer with a learning rate of 0.001, β_1 of 0.5, β_2 of 0.9, and weight decay of 1×10^{-7}. The hyperparameters utilized follow those used in previous work [29,44], where extensive ablation studies were conducted, with the exception of the batch size which was selected to be the largest value that could be used on our training hardware. Training progression was tracked using a mismatch term equal to the mean Frobenius norm of the difference between the true fire arrival time $\boldsymbol{\tau}$ from the training set and generated fire arrival time samples $\boldsymbol{\tau}^g$, for a given set of inputs $(\bar{\boldsymbol{\tau}}, \boldsymbol{h})$. The final model state selected for evaluation was determined based on performance on separate validation data, choosing the training checkpoint which provided the best performance across the cases examined.

2.4. Estimating Ignition Times from GOES

When applying the cWGAN-based approach to real fires, ignition time estimates are gathered from the Geostationary Operational Environmental Satellite (GOES) Fire/Hot Spot Characterization product. The GOES fire product uses both visible and infrared spectral bands to locate actively burning fires with a spatial resolution of 2–4 km depending on the location being measured (coarser for locations farther from the nadir) and a temporal resolution of 5 min [22]. The GOES fire product provides a geolocated fire mask which categorizes each pixel as either fire with an associated confidence level or as another non-fire category. GOES data has been collected using AWS, with GOES-West being utilized for the fires examined here, which occurred in the Western US, though data from GOES-East may also be utilized. To gather ignition time estimates from the GOES fire product, the following steps are followed:

1. GOES fire mask data from the Fire/Hot Spot Characterization product is collected for a time period of +/− 1 h around the approximate fire start time.
2. A domain of size 12.8 km × 12.8 km is centered around the fire of interest.
3. The portion of the collected fire masks intersecting the chosen domain is identified based on latitude and longitude coordinates of fire mask pixels.
4. Fire masks containing an active fire pixel (excluding low or nominal confidence pixels) in the selected domain are determined.

5. The time of the earliest detection in the domain of interest is taken as the estimated ignition time.
6. In the case that no GOES fire detection is found, additional data is collected for a larger time period around the approximate start time and the above steps are repeated until a detection is found.

Following this set of steps, an ignition time estimate is obtained for a fire of interest, which is subsequently utilized to construct the measurement input to our network.

2.5. Constructing Measurements from VIIRS Active Fire Data

Active fire measurements are obtained from the Collection 2 Visible Infrared Imaging Radiometer Suite (VIIRS) 375 m Active Fire (AF) Product. These measurements are obtained via the polar-oribiting Suomi-NPP and NOAA-20 satellites, which provide data products VNP14IMG and VJ114IMG, respectively. The VIIRS AF product uses infrared imaging to determine the location of actively burning fires during a 6 min measurement period on a grid with 375 m resolution [21]. The final data product provides the latitude and longitude coordinates for each AF detection pixel, along with a corresponding confidence level of high, nominal, or low. Additionally, a categorical fire mask is provided which supplies classifications for each 375 m pixel, including non-fire pixels, on a unique geolocated grid for each measurement file, where the location and orientation of the grid is dictated by satellite position during the measurement. Measurements are provided anywhere from 2 to 4 times a day per satellite, depending on the geographic location of interest relative to the orbit of the satellites. VIIRS AF data is collected from NASA LAADS DAAC based on the time period of interest and approximate fire location.

Following collection of VIIRS AF data for a fire of interest, a geolocated discretized grid of size 34×34 with 375 m resolution centered over the fire of interest and co-located with the domain used for determining ignition times from GOES is created. Based on the latitude and longitude coordinates of high-confidence AF detections, pixels in the discretized grid are assigned arrival time values based on the VIIRS measurement times and the ignition time estimate obtained from GOES (i.e., they are assigned a value equal to the VIIRS measurement time minus the GOES-derived ignition time, to provide an approximate arrival time value). Thereafter, unassigned pixels in the discretized grid are assigned a value of 48 h, following what is performed for measurements in the training set. The measurement image is then resampled to the 25 m resolution used by the cWGAN, and padded to return to a total domain size of 12.8 km $\times$ 12.8 km.

3. Results

In the following section, we apply the cWGAN-based approach for fire arrival time reconstruction to wildfires taking place between 2019 and 2022 in California and Washington. We first discuss the wildfire incidents considered, after which we present the GOES-derived ignition times for these cases and the VIIRS AF data utilized. Following this, the fire arrival time reconstructions are presented and validated using high-resolution fire perimeters, and prediction uncertainty is discussed. Lastly, the influence of terrain conditioning on the fire arrival time reconstructions is examined.

3.1. Wildfire Incidents

The selection of wildfire incidents for validation of the cWGAN-based approach for fire arrival time estimation is based on two primary factors. First, the extent of the wildfire after the initial 48 h of burning should be contained within a domain of size 12.8 km by 12.8 km, as prescribed by the problem setup and training data utilized. Second, to allow for the validation of generated fire arrival time predictions, fire perimeters must be available

within the first 48 h of the fire. For this purpose high resolution infrared (IR) fire perimeters from the National Infrared Operations (NIROPS) program are utilized. These perimeters are collected via IR sensors mounted aboard aircraft which fly over wildfires and map their extent at the time of measurement with resolutions ranging from 0.46 to 3.5 m/pixel [47,48]. Following these considerations, five wildfires were selected which exhibit a variety of sizes, terrain conditions, and AF detection densities. These are the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires. Additional information about these fires is provided in Table 1.

Table 1. Wildfire incidents utilized for model validation, along with their ignition date, GOES-derived ignition time, date of first IR perimeter measurement, and approximate North–South and East–West extent. All reported dates and times are in UTC.

Fire	Ignition Date	GOES Ignition Time	IR Measurement Date	IR N-S Extent	IR E-W Extent
Bobcat	6 September 2020	19:16	8 September 2020	11.9 km	7.8 km
Tennant	28 June 2021	23:21	30 June 2021	11.4 km	7.7 km
Oak	22 July 2022	21:26	24 July 2022	10.2 km	9.7 km
Barnes	7 September 2022	23:51	9 September 2022	3.6 km	7.9 km
Williams Flats	2 August 2019	14:46	4 August 2019	6.5 km	11.2 km

For the fires selected, ignition time estimates are obtained from the GOES Fire/Hot Spot Characterization product and are used in conjunction with VIIRS AF data to construct the measurement inputs to be utilized. Terrain data is additionally collected for the domain of interest. Figure 3 depicts the final measurements constructed for the fires examined here, along with their corresponding terrain height maps.

3.2. Fire Arrival Time Predictions

The arrival time measurements and terrain height maps for the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires are used as input to the trained generator of the cWGAN model. For each measurement–terrain pair, 500 corresponding fire arrival time map instances $\boldsymbol{\tau}^g$ are generated from the learned conditional distribution $P^g_{\boldsymbol{T}|\bar{\boldsymbol{T}},\boldsymbol{H}}$ by sampling the latent vector $\boldsymbol{z}$ from its distribution $P_{\boldsymbol{Z}}$ and passing these through the generator. Generated samples are used to compute the pixel-wise mean prediction for each fire following Equation (4), providing the best estimate of fire progression. The pixel-wise standard deviation is also computed across the generated samples to provide an estimate of prediction uncertainty. Figure 4 shows the resulting mean fire arrival time predictions and the standard deviation of generated fire arrival times for the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires.

Using geolocation information for predicted fire arrival times, we can also visualize fire progression relative to geographical features visible through satellite imagery. Figure 5 presents contour maps of the mean fire arrival time predictions generated for the fires of interest, displayed in Google Earth, with contours at 4 h intervals. In these images, IR perimeters collected from NIROPS are also provided (in blue) as a reference.

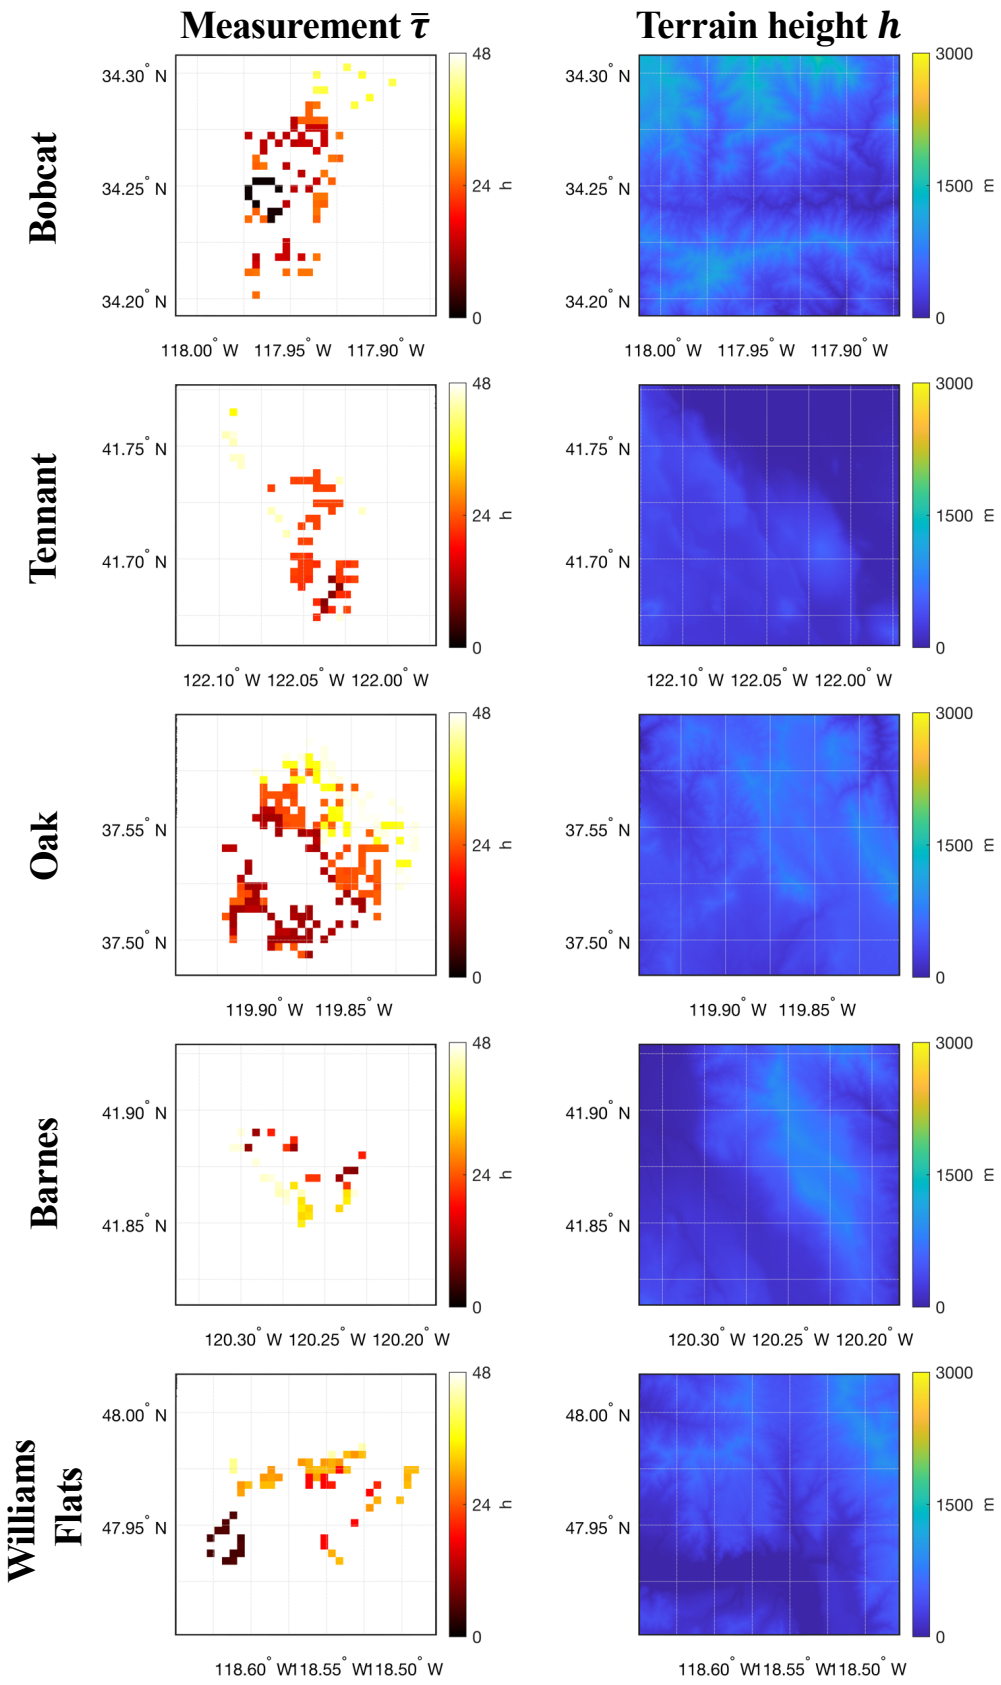


Figure 3. Inputs to the cWGAN model for the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires. The inputs include measurements based on VIIRS high-confidence detections with arrival times calculated from GOES ignition times (**left**) and terrain height maps (**right**).

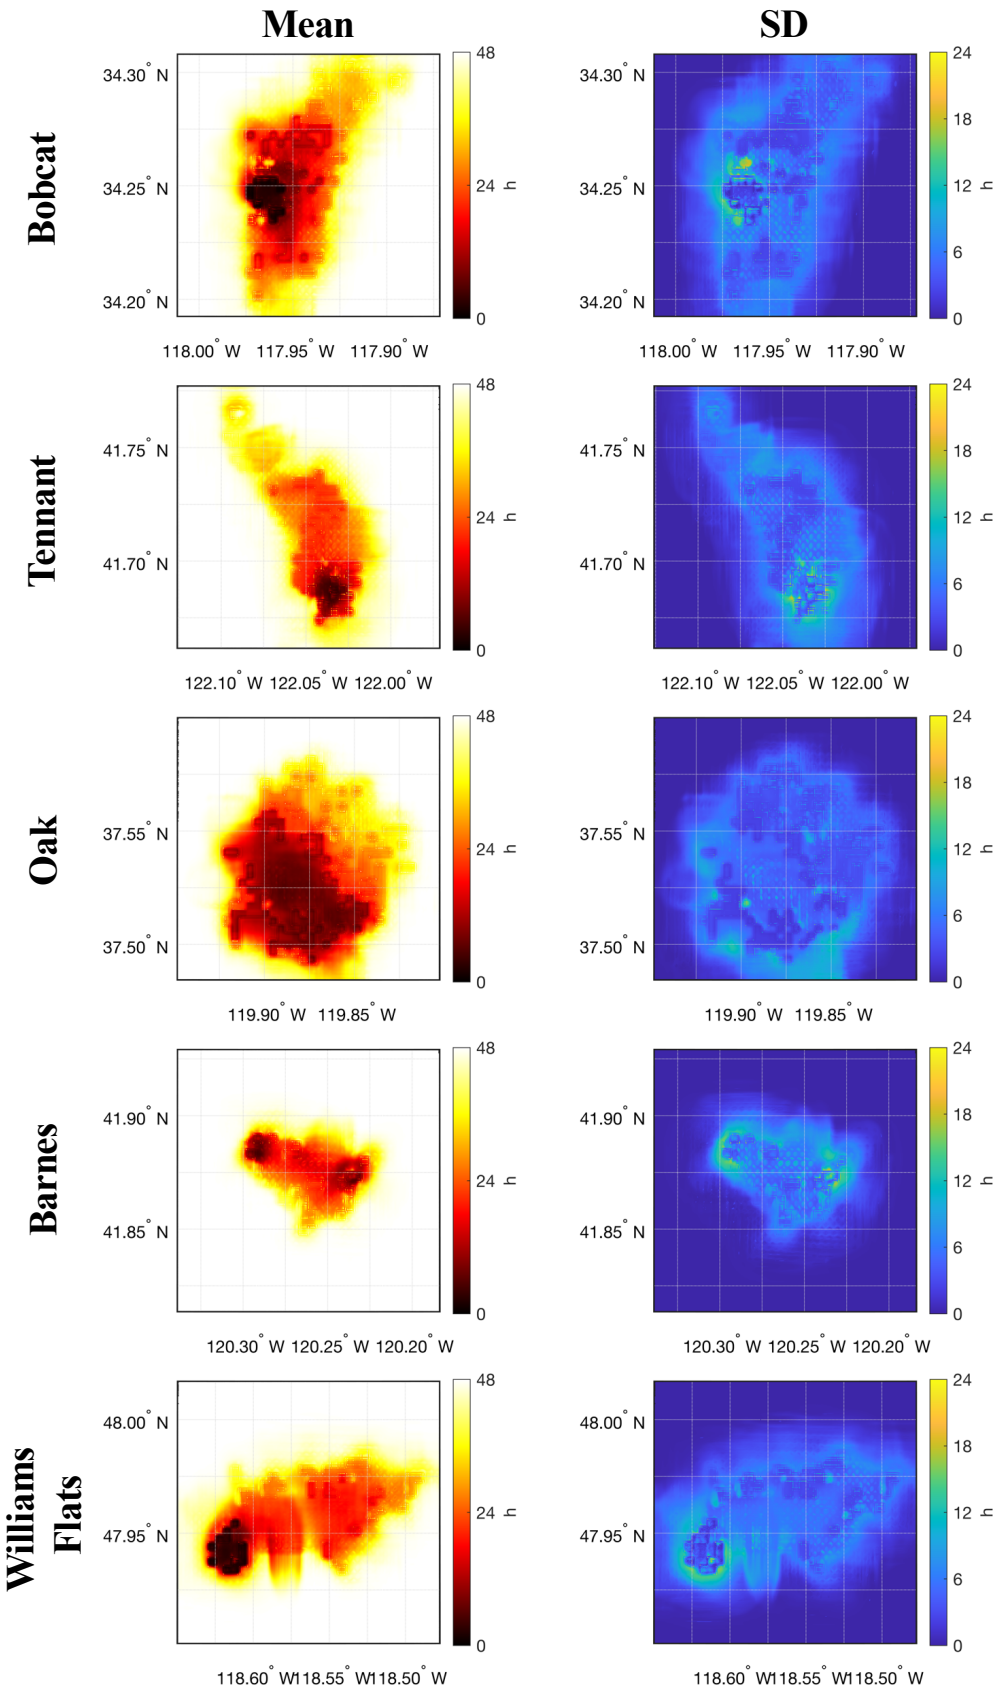


Figure 4. Mean fire arrival time predictions (**left** column) and standard deviation of fire arrival time predictions (**right** column) computed using 500 generated fire arrival time instances for the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires.

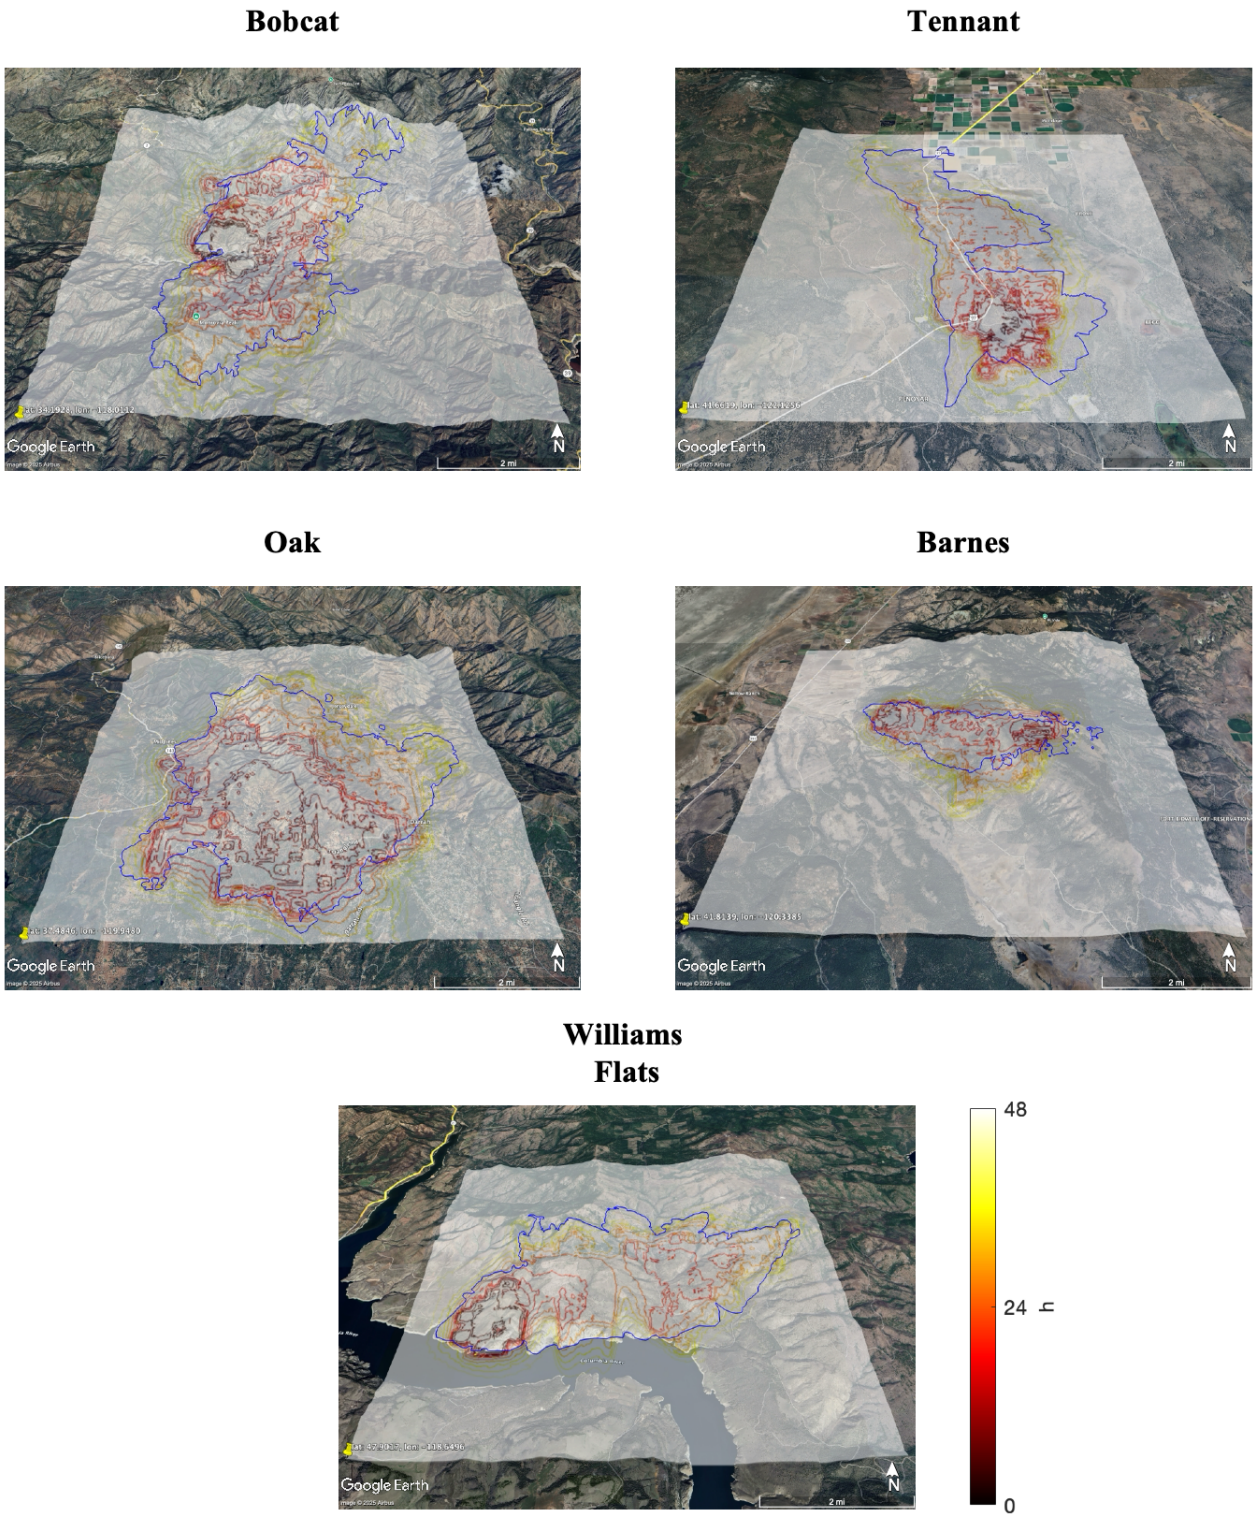


Figure 5. Contour maps of mean fire arrival time predictions for the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires displayed in Google Earth, with contours at 4 h intervals. Measured IR perimeters are indicated in blue for reference.

3.3. Prediction Validation

To validate predicted fire arrival time maps, we compare estimated fire perimeters against available IR fire perimeters from NIROPS. Predicted fire perimeters are determined from the mean fire arrival time maps by taking contours corresponding to the times that the IR perimeters were measured. Based on the geolocation information for the fire arrival time maps, estimated perimeters are geolocated, making them directly comparable

to measured IR perimeters. Based on the intersection of the estimated and measured perimeters three areas are determined, (1) area of agreement, (2) false negative area, and (3) false positive area, where areas of agreement account for areas where both the predicted and measured perimeters indicate fire presence, false negative areas account for regions where the measured perimeter indicates fire presence but the predicted perimeter does not, and false positive areas account for regions where fire is predicted but not present in the measured perimeter. Figure 6 displays predicted perimeters for the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires along with measured IR perimeters, with the three areas of interest indicated.

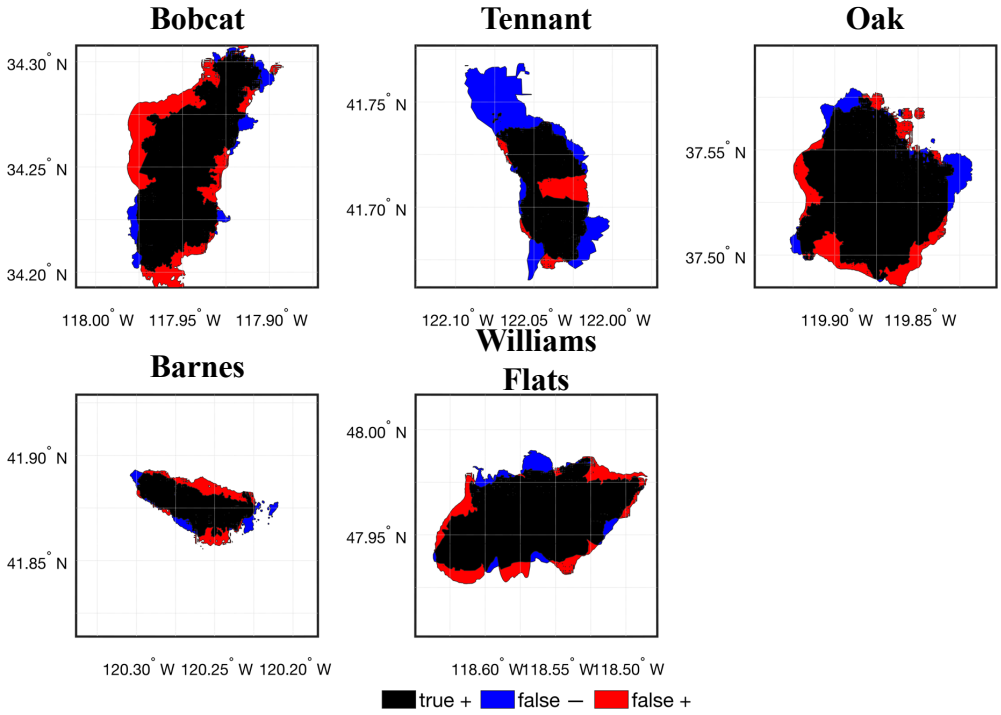


Figure 6. Comparison of predicted fire perimeters against measured high-resolution IR fire perimeters from NIROPS, with areas of agreement colored in black, false negative areas colored in blue, and false positive areas colored in red.

We compare predicted and measured perimeters quantitatively by evaluating the Sørensen–Dice coefficient (SC), the probability of detection (POD), and the false alarm ratio (FAR). SC is a measure of how well the two perimeters agree, POD is a measure of how well the prediction captures true fire growth, and FAR is a measure of false growth in predictions. These metrics are defined as

$$SC = \frac{2A}{2A + B + C} \qquad POD = \frac{A}{A + B} \qquad FAR = \frac{C}{A + C} \tag{5}$$

where A is the area of agreement between the two perimeters, B is the false negative area, and C is the false positive area. For all three metrics values range between 0 and 1, with the optimal value for the SC and POD being 1 and the optimal value for the FAR being 0.

Table 2 contains the SC, POD, and FAR values for the fires evaluated here. We observe that for all cases except the Tennant fire, the SC values are above 0.80, which indicates that the predicted fire extent agrees well with the measured IR perimeters. Additionally, POD values for Bobcat, Oak, and Williams Flats are above 0.90, indicating that, in these cases, the model does an excellent job of capturing the total fire growth present in the measured perimeters. In all cases, the FAR values are less than 0.30, with values for three of the cases additionally being below 0.20, indicating a low tendency to falsely predict fire where fire is

not present. For the Tennant fire, we observe especially low SC and POD values resulting from the particularly sparse measurement and unique spread pattern observed. Overall, these metrics indicate that, in most cases, the model reasonably captures wildfire extent at a given time, while balancing trade-offs between overpredicting and underpredicting the extent.

Table 2. Sørensen–Dice coefficient (SC), probability of detection (POD), and false alarm ratio (FAR) values for predicted fire perimeters relative to high-resolution measured IR perimeters from NIROPS for the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires. SC, POD, and FAR values achieved by the cWGAN-based approach of Shaddy et al. [29] are additionally included for the Bobcat, Tennant, and Oak fires.

	New Method			Shaddy et al. [29]		
Fires	**SC**	**POD**	**FAR**	**SC**	**POD**	**FAR**
Bobcat	**0.82**	0.91	**0.26**	0.80	**0.97**	0.32
Tennant	0.67	0.57	**0.17**	**0.78**	**0.78**	0.21
Oak	**0.87**	0.90	**0.16**	0.84	**0.97**	0.26
Barnes	0.81	0.87	0.25	-	-	-
Williams Flats	0.88	0.92	0.16	-	-	-

Bold values indicate the better score in each metric between the two approaches.

To further evaluate performance, we compare SC, POD, and FAR values to those obtained in Shaddy et al. [29] for the Bobcat, Tennant, and Oak fires (included in Table 2). For the Bobcat, Tennant, and Oak fires which are assessed in Shaddy et al. [29] and evaluated against the same IR perimeters as used here, we find that the new approach exhibits better performance on the Bobcat and Oak fires when measured using SC. Comparing POD and FAR values between the two approaches, we find that the approach of Shaddy et al. [29] scores better in POD across all cases, though with the approach presented here achieving better scores for FAR. We note that there are a number of differences between the approach developed here and that of Shaddy et al. [29], including differences in both problem setup and the hardware/environment used for training and inference. For the problem setup, notable differences include an increased resolution of 25 m in the present work, compared to the 60 m resolution used in [29]. Additionally, in the present work, fire arrival time estimates are conditioned on ignition time information from GOES, which was not utilized in [29], and further uses data from both the NOAA-20 and Suomi-NPP satellites, as opposed to relying solely on data from Suomi-NPP. Regarding the hardware and environment used, differences include the GPU hardware used for training and inference, the training batch size, the number of epochs trained, and the number of training points utilized.

3.4. Evaluation of Terrain Influence on Arrival Time Reconstruction

To evaluate the influence of terrain height conditioning on fire arrival time estimates, additional tests are undertaken. For a given measurement, fire arrival time predictions are conditioned on both the correct terrain height information and a constant zero-valued non-varying terrain height. For the 200 measurement samples from the validation dataset, we generate 500 fire arrival time maps both for the correct terrain height data and for flat terrain. Thereafter, we compute the mean arrival time predictions for the two cases. Finally, we compute the pixel-wise difference in arrival times by subtracting the mean fire arrival time map for the case with the correct terrain information from the mean fire arrival time map for the case with flat terrain. This is done for each of the 200 samples considered from the validation set, and three representative cases are shown in Figure 7.

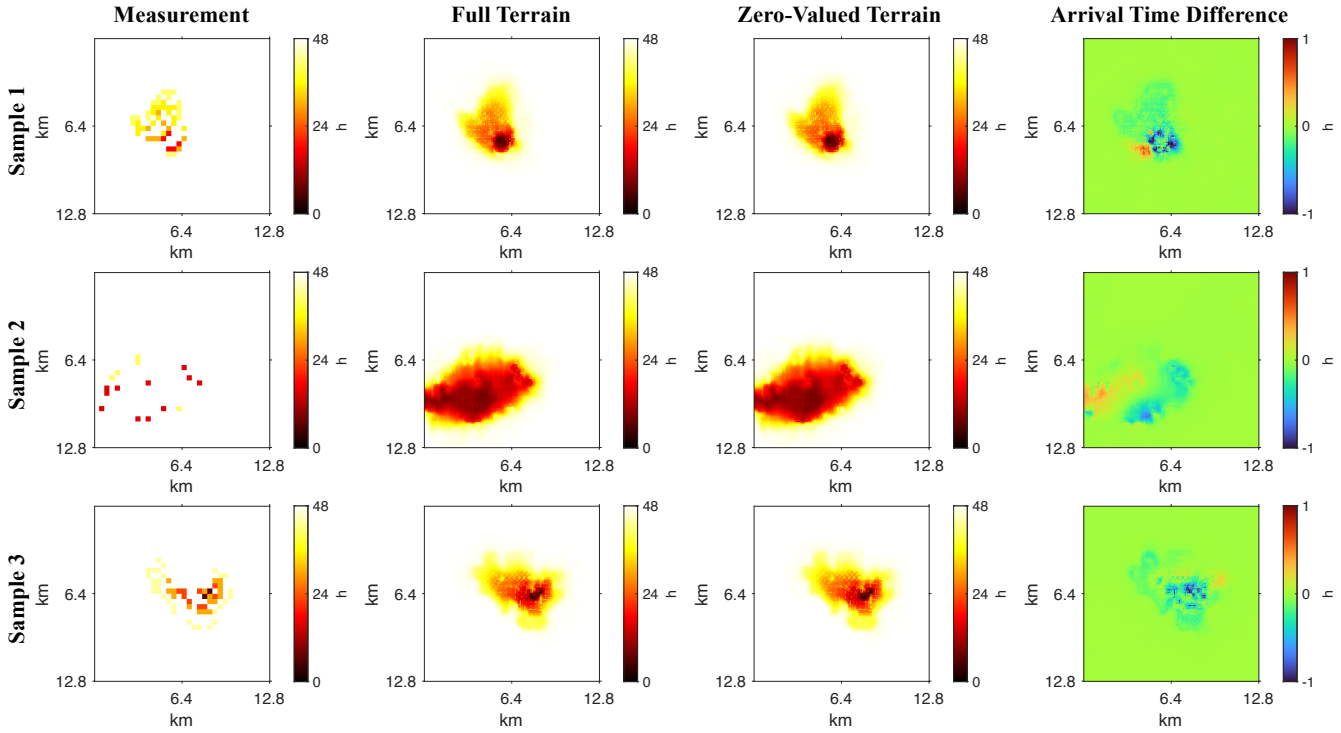


Figure 7. Difference between mean predicted fire arrival time maps conditioned on full terrain maps and zero-valued non-varying terrain maps for three sample cases. The conditioning measurement is shown in the first column, followed by the mean fire arrival time predictions produced for the full terrain maps, the mean fire arrival time predictions produced for the zero-valued terrain maps, and the difference between the two mean arrival time maps.

From the 200 samples assess, we further compute a histogram of the pixel-wise differences, which are plotted in Figure 8. Unburned regions, which are not of interest, are excluded from this calculation by removing any pixels with arrival time values exceeding 47 h in both mean predictions. From Figure 8, we observe that more than 97% of pixels have a difference in predicted fire arrival time values of less than 30 min.

To further examine the influence of terrain we consider a case where the measurement is not sufficiently sampled for the entire fire event, namely the Oak fire. In the case of the Oak fire, no measurements are available until more that 11 h after ignition, resulting in a region in the center of the fire with no detections. The central region of the Oak fire is additionally bounded by later-time detections which were observed at the same measurement time, making it further challenging to determine how the fire may have progressed through this region. We prescribe two artificial terrain maps for Oak, one which slopes upward from the southwest corner to the northeast corner and one with an opposite gradient, providing a terrain slope which is roughly parallel to the primary apparent growth direction. To assess the influence of terrain in this region, we aim to look for differences in the predicted ignition location for the two prescribed terrain maps. Fire is known to spread faster uphill and slower downhill due to effects such as preheating of fuels, so it would be expected that within this region the predicted ignition location would move downhill towards detections on one side of this gap in measurements, thus resulting in a fire which progresses slower downhill and faster uphill before ultimately reaching the detections on either side of the prescribed slope at the same time.

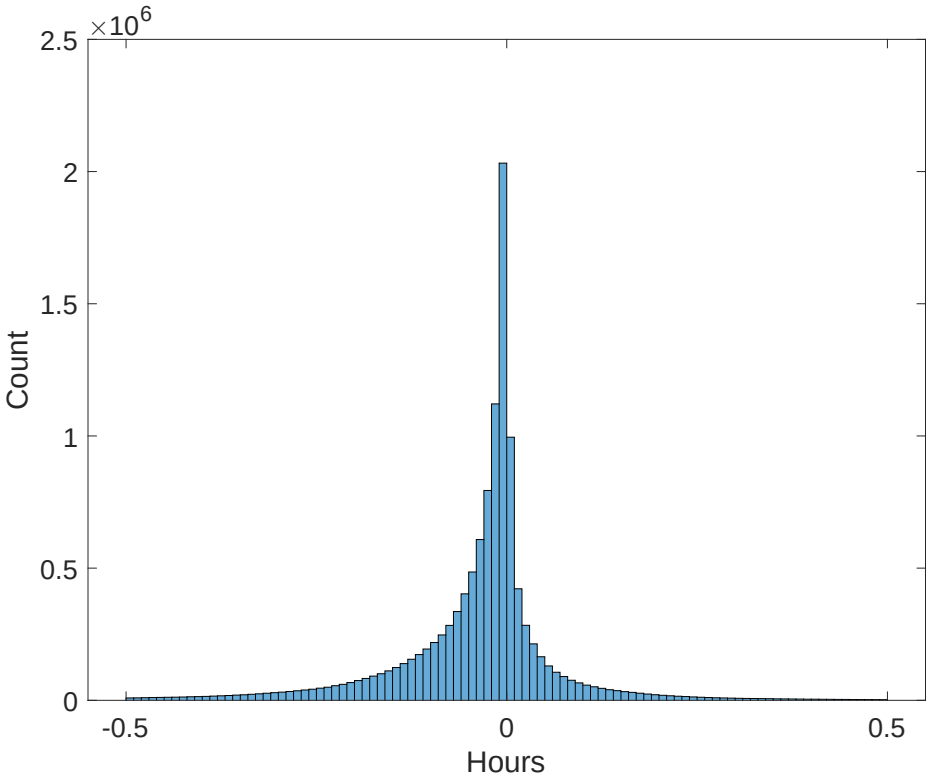


Figure 8. Histogram of pixel-wise difference in mean predicted fire arrival time maps conditioned on correct terrain and their counterparts conditioned on a flat terrain. The data was generated for 200 wildfires selected from a validation set.

Figure 9 presents the measurement and terrain maps for this case in panels (a) and (b), along with the mean fire arrival time estimate with the predicted ignition location indicated. We observe that for the prescribed terrain maps, the predicted ignition location shifts in the downhill direction, towards measurements on the downhill side of the sparsely measured region, matching what is to be expected.

To illustrate the dominant effect of measurements relative to terrain data, we add artificial early-time detections at 4 h after ignition to the center of the measurement for the Oak fire and repeat this test. Figure 9 displays the results from this test in panels (c) and (d), which indicate that there is no longer a significant shift in the predicted ignition location due to the changing terrain.

A final test is performed utilizing the measurement for the Williams Flats fire, which contains a cluster of early time detections spread over a sizable region in the Western portion of the fire which were measured roughly at the same time. Again, we alter the terrain map, prescribing both an uphill slope from West to East and an opposite downhill slope, and estimate fire arrival times. Figure 10 shows the resulting mean fire arrival time predictions and estimated ignition locations. In this case, we observed a shifting of the predicted ignition location within the cluster of early-time detections towards the downhill direction.

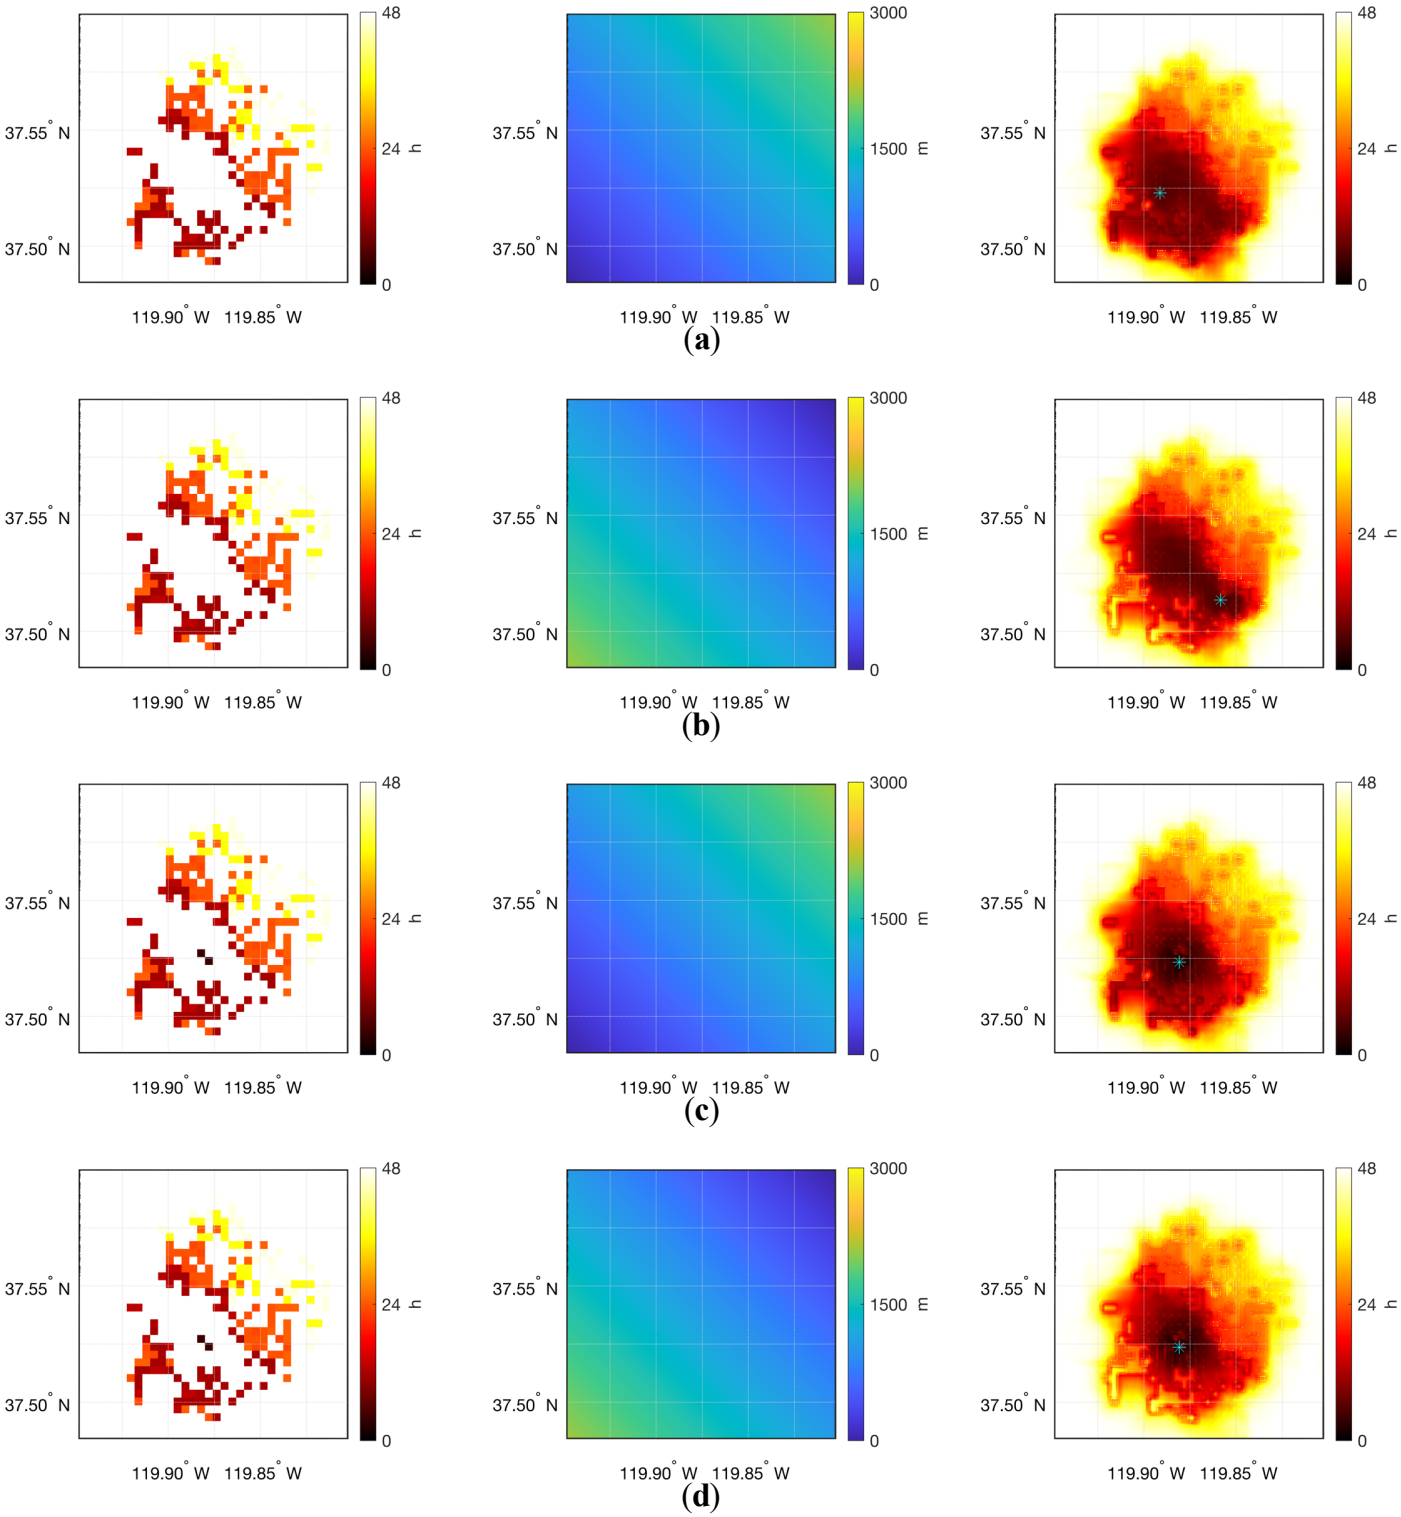


Figure 9. Oak fire measurements and prescribed terrain maps with gradient across sparsely measured region, along with mean predicted fire arrival time maps with predicted ignition locations indicated using a blue star, for (**a**) true measurement with an upward slope from southwest to northeast, (**b**) true measurement with a downward slope from southwest to northeast, (**c**) measurement with artificial early-time detections and an upward slope from southwest to northeast, and (**d**) measurement with artificial early-time detections and a downward slope from southwest to northeast.

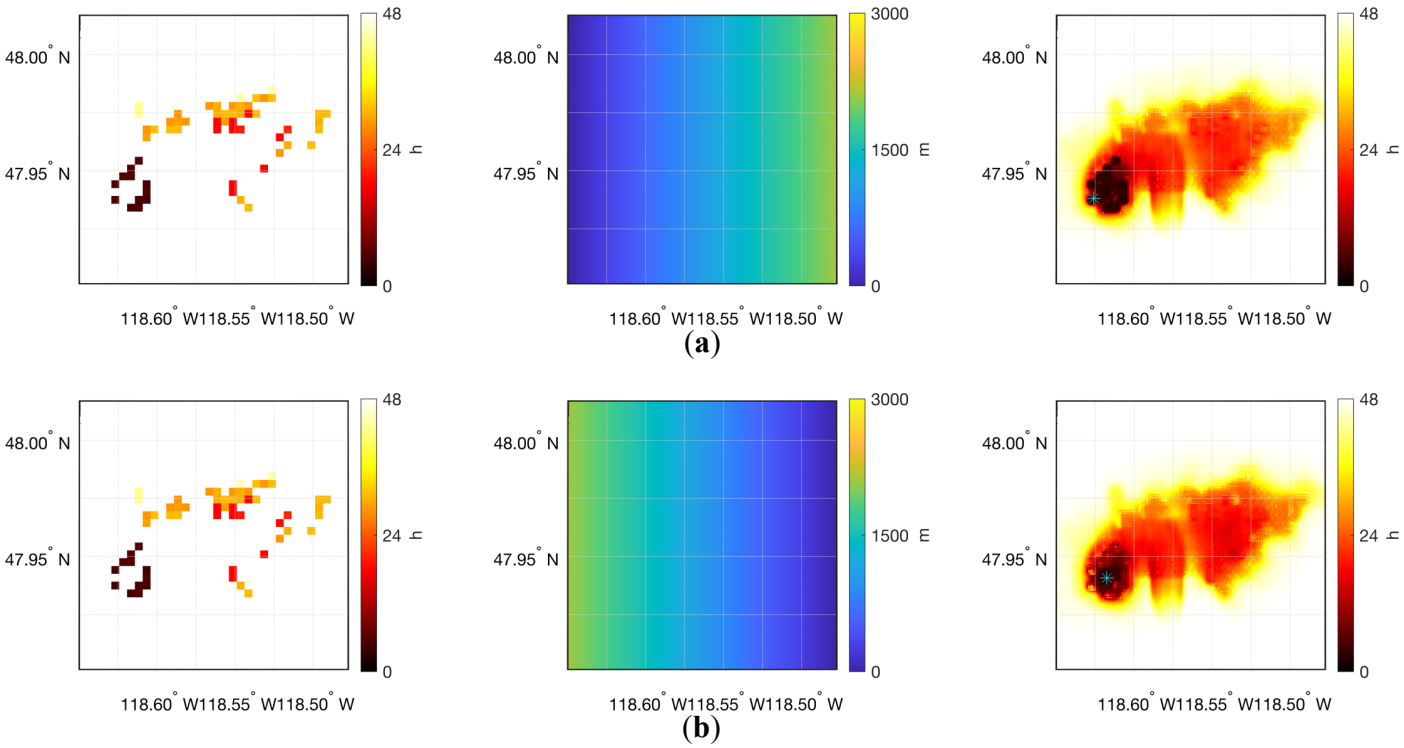


Figure 10. Measurement for the Williams Flats fire and prescribed terrain maps, along with the mean predicted fire arrival time maps with ignition locations indicated using a blue star, for (**a**) an upward-sloping terrain from west to east and (**b**) a downward-sloping terrain from west to east.

4. Discussion

Examining the mean and standard deviation of fire arrival time estimates produced for the Bobcat, Tennant, Oak, Barnes, and Williams Flats fires presented in Figure 4, we see that the mean fire arrival time smoothly interpolates the AF detections provided as input, resulting in arrival time maps which may be used to generate a progression of fire perimeters consistent with the measured behavior. We also observe that even in cases where there is an absence of measurements for a large potion of the fire, such as the center of the Oak fire or the southern portion of the Williams Flats fire, the predicted fire progression captures fire growth through these regions. Additionally, we observe a higher standard deviation around the periphery of the fires, indicating that there is higher uncertainty about the final extent of the fires at the end of the initial 48 h. We note that arrival time predictions in locations corresponding to measurements are strongly influenced by provided measurement times, resulting in a lower standard deviation in these regions. We see also that regions corresponding to a sharp gradient of arrival times in the measurement, such as the western portions of the Bobcat and Williams Flats fires, correspond to a higher uncertainty, as evidenced by the elevated standard deviation. This increased standard deviation appears to stem from the abrupt difference in measured arrival times between neighboring detections, which implies an abrupt and non-smooth fire progression.

Further examining Figure 5, which presents geolocated contour maps of the mean fire arrival times displayed in Google Earth, we see an example of how contextualized visualizations can provide a sense for how the fire front expanded over the course of each fire, relative to visual characteristics of the landscape. These images display a sequence of predicted perimeters, from which intuition can be gathered about the behavior of each fire as it relates to terrain, vegetation, and other visible geographical features. Visualizations

such as these may also prove useful for fire managers, where obtaining a measure of a fire's current extent relative to the landscape could help with planning fire response.

Looking at the agreement between the mean predicted perimeters provided by the cWGAN approach and the IR perimeters from NIROPS, presented in Figure 6, we observe that in most cases, the cWGAN-based approach appears to capture the overall shape and extent of the wildfires at the time of IR perimeter measurement, with the exception of some finer spread patterns around the periphery which are not fully captured. However, as noted, the standard deviation tends to be elevated in these regions around the periphery where there is higher error, indicating higher uncertainty. We believe that these fine-scale details likely stem from small-scale effects arising from interactions between weather, terrain, and fuel during fire spread, which the cWGAN model has no knowledge beyond the limited information embedded in the measurement and terrain inputs. In the presence of higher spatial- and temporal-resolution measurement data it may be possible to capture more of these details; however, even in this case, the stochasticity of fire spread remains as a limiting factor. In examining comparisons between the predicted and measured perimeters further, we see that in most cases, false positive area and false negative area tend to be similar in magnitude, indicating that the model does not display an obvious bias.

In the case of the Tennant fire, the model under-predicts growth in the north and incorrectly predicts burning in the central region. We believe the northern false negative region arises from the large temporal gap in measurements, which jumped from roughly 22 h after ignition to 45 h after ignition, with the exception of a single detection around 36 h in the northern-most position for this set of measurements. This temporal gap means that very little information about fire progression during the second day of growth was provided, resulting in the northern progression being predicted to occur later than what is indicated by the IR perimeter. From Figure 4 we do observe that the standard deviation in this region is slightly elevated relative to the center-most correctly predicted region, indicating higher uncertainty. Regarding the false positive growth in the central region of the fire, we believe that the growth pattern for the Tennant fire is somewhat unique relative to other fires examined here, for instance the Williams Flats fire had a similar region with no detections; however, in that case, fire did spread through the region and was correctly predicted by the cWGAN approach. It may also be the case that the fuel (vegetation) present played an important role in this unique growth pattern, which our model currently is not conditioned on, though this conditioning may be a beneficial addition in future work.

Comparing the performance of the approach developed here for fire arrival time inference and that presented in Shaddy et al. [29], our results indicate that the method developed here has a lower tendency to overpredict wildfire extent relative to IR perimeters, which results in a slightly lower POD value, but with fewer false positives. SC remains the most effective way to measure overall agreement with measured IR perimeters, however, for which the approach presented here produces better results in two out of three cases. To explain the difference in performance for the Tennant fire, for which our new approach falls short of the previous work, we note that Shaddy et al. [29] utilized nominal confidence VIIRS detections, in addition to the high-confidence detections used here, which filled the large temporal gap present in the high-confidence data which we have discussed. Due to the similar issues of sparsity in high-confidence data, we do not evaluate the Mineral fire, which was examined in Shaddy et al. [29].

Further examining the overall performance of the fire arrival time inference approach presented here, we find that our method achieves an average SC of 0.81 across the five cases used for validation, relative to NIROPS perimeters which are treated as ground-truth. When the Tennant fire is excluded, the average SC coefficient jumps to 0.85, demonstrating the importance of having sufficient measurement data. Comparatively, we find that for the

cWGAN approach presented in [29], an average SC of 0.81 between predicted perimeters and ground-truth NIROPS perimeters is reported across the four fires which were used for validation in that work (Bobcat, Tennant, Oak, and Mineral). [29] also presented results for the SVM approach of [28] on the same validation fires, reporting an average SC of 0.79 across the four validation cases examined. Looking at the performance of the approach presented in McCarthy et al. [35], which utilized a CNN to downscale GOES data to achieve fire progression estimates with 375 m spatial resolution and 5 min temporal resolution, authors report intersection over union (IoU) values achieved by their approach for five fires, relative to perimeters from GeoMAC. These IoU scores are converted to SC values, finding an average of 0.72. Overall, we believe that these results indicate that the cWGAN-based approach developed here has achieved a respectable improvement over previous work, though with some dependence on the availability of measurement data for a given fire, while also providing advancements including capabilities for uncertainty quantification. We note that a more comprehensive comparison requiring the collection of codes and data for competitive models and assessment fn the same fires examined here is beyond the scope of this paper.

To assess the influence of terrain conditioning on generated fire arrival time maps, we have looked at the influence terrain has on predicted fire arrival time maps by conditioning on both the correct terrain maps and non-varying terrain maps for a number of cases from a testing dataset. From Figure 7, we saw that discrepancies between the mean arrival time predictions conditioned on proper terrain and non-varying terrain tend to occur around the periphery of a fire, though the differences are largest in regions where measurements are sparse or uninformative. In these situations, the measurement becomes less informative, leaving the cWGAN to rely on any information about the fire progression which may be gleaned from the terrain data. This further suggests that the inference is most strongly influenced by the satellite measurements, while terrain appears to play a secondary role.

The histogram of arrival time differences arising from conditioning on the full terrain information and non-varying terrain data, presented in Figure 8, further suggests that predictions are most strongly influenced by arrival time measurement data, and once this data has been assimilated, the effect of the terrain map becomes less pronounced. It is important to note that the measurement itself contains information about how terrain has influenced fire progression; thus, given a sufficiently sampled measurement, terrain data becomes less informative, but in cases where the measurement is not sufficiently sampled, the terrain plays a critical role, as is shown in Figure 7.

The assessment of the terrain influence on fire arrival time estimates and predicted ignition locations presented in Figure 9 and Figure 10 was further able to demonstrate that in cases where many detections made at the same time are available for a large region, the terrain data can strongly influence the spread of the fire within that region, in addition to playing an important role in measurement sparse regions. This is to say that when a set of detections is insufficient to resolve fire progression within a particular region, the knowledge of the terrain can strongly influence the fire progression estimate.

5. Conclusions

In this work, we have developed an approach for estimating wildfire progression from satellite measurements and terrain height data using a conditional Wasserstein Generative Adversarial Network (cWGAN) trained on simulations of historic wildfires. The trained cWGAN generates samples from the conditional distribution of fire arrival times given satellite active wildfire detections, augmented with ignition time estimates from geostationary satellites, and terrain height data. Owing to the probabilistic framework of the approach, uncertainties can also be quantified by computing the standard deviation of

generated samples, and further, a best guess prediction is found from the mean of generated samples. Once obtained, reconstructed fire arrival time maps may be used to initialize forecasts from coupled atmosphere–wildfire models or for situational awareness.

Training data was generated from WRF-SFIRE wildfire spread simulations of the initial 48 h of 140 wildfires occurring in CONUS during 2023. Artificial satellite measurement data for training is obtained by applying an approximate observation operator to simulated fire arrival time maps, producing measurements that are consistent with true satellite data. By training with solutions from WRF-SFIRE, the wildfire-spread physics of this model are incorporated into fire arrival time estimates produced by the trained cWGAN.

When applied to real fires, active fire detections from the polar-orbiting VIIRS sensor were utilized, along with ignition time estimates obtained from GOES and terrain height data. With satellite measurements and terrain data collected and preprocessed, the method could be applied to generate a single fire arrival time estimate in under 0.1 s, and 500 samples could be produced and resulting statistics computed in under 1 min. The trained cWGAN was applied to wildfires occurring in California and Washington, and the resulting mean predictions were evaluated against high-resolution IR wildfire perimeters provided through the NIROPS program. Predicted wildfire perimeters were compared to IR perimeters, and a mean Sørensen–Dice coefficient of 0.81 was found across five wildfire instances (0.85 when excluding the Tennant fire), along with an average probability of detection of 0.83, and an average false alarm ratio of 0.20. Of the five fires examined, predictions for the Tennant fire exhibited the largest discrepancy relative to the measured IR perimeter, appearing to result from insufficient AF detections which led to a temporal gap of roughly 20 h in measurement availability. Results were additionally compared to those presented in Shaddy et al. [29] for the Bobcat, Tennant, and Oak fires, finding that the approach presented here led to improved performance for the Bobcat and Oak fires when evaluated using the Sørensen–Dice coefficient and additionally had a lower false positive rate across all cases, as indicated by the false alarm ratio. The influence of terrain height on fire arrival time predictions was also explored, revealing that terrain had the greatest impact in regions where measurements provided limited information about fire progression. The increased influence of terrain on the fire arrival time inference was demonstrated for regions of a fire where measurements were particularly sparse and for regions where many detections spread over a broad area were observed at the same time.

We note that the approach developed here has several shortcomings, which may be improved upon in future work. These include the domain size considered, which prohibits the use of this model for evaluating particularly large fires, and the limited trajectory length for fire arrival time reconstructions. Both points may be addressed through modification of the training data. Namely, to address the limited domain size, the training data may be modified by reducing the resolution or increasing the grid size, allowing the current set of WRF-SFIRE results to be utilized to obtain training data for a larger domain. However, increases in grid size come at the cost of a larger cWGAN architecture, leading to more costly training and evaluation. To modify the reconstruction time frame to accommodate longer fire arrival time trajectories, the WRF-SFIRE simulations used for training may be extended to be arbitrarily long, though with a reduction in their accuracy for longer time horizons. In this case, the use of sequential data assimilation may be more appropriate. We further note that the framework developed here may be applied to any wildfire progression data, following the same set of steps described here to generate training data and apply the model, making the utility of this approach valid regardless of the accuracy of the particular wildfire spread solutions used to demonstrate the approach here.

Another limitation, stemming from the cWGAN model used here, is the “imprinting” of the measurement image on predicted fire arrival time maps observed in some cases. This

results from the U-Net architecture for the generator, which includes skip connections and directly transforms measurement and terrain maps into fire arrival time predictions. This may be overcome through use of another type of conditional generative model, such as a diffusion model, wherein Gaussian noise is sequentially transformed into the target result, which would likely address the problem of imprinting [49,50]. We additionally note that due to the regression loss used, the training of these models is also often easier, and it is typically more straightforward to asses when training is completed based on their loss curve. However, the sampling stage requires multiple steps (on the order of 100), which makes sample evaluation more expensive.

Future extensions of this work will focus primarily on inclusion of additional conditioning variables such as fuel type, fuel moisture, and weather conditions. We believe that, by including additional physical variables relevant to wildfire spread into the inference, the model will learn to capture more complex physical relationships and subsequently incorporate these relations into fire arrival time estimates.

Author Contributions: Conceptualization, B.S., A.D. and A.A.O.; methodology, B.S., A.D. and A.A.O.; software, B.S.; validation, B.S.; formal analysis, B.S.; investigation, B.S. and H.Q.; resources, K.H. and A.A.O.; data curation, B.S., B.B., A.F. and J.H.; writing—original draft preparation, B.S.; writing—review and editing, all coauthors; visualization, B.S.; supervision, A.A.O., K.H., J.M., D.V.M. and A.K.; project administration, K.H. and A.A.O.; funding acquisition, K.H., A.K., J.M., D.V.M. and A.A.O. All authors have read and agreed to the published version of the manuscript.

Funding: The NOAA Bipartisan Infrastructure Law project NA22OAR4050672I supported the creation of the forecasts. The NASA Disasters project 80NSSC19K1091 supported the high-performance computing for running the forecasts. The NASA FireSense project 80NSSC23K1344 supported development and evaluation of the machine learning approach. JM was funded by NASA grants 80NSSC22K1405, 80NSSC23K1118, and 80NSSC23K1344. AK acknowledges partial support from NASA 80NSSC23K1118 and NSF-DEB 2039525. The APC was funded by NASA grant 80NSSC23K1344. AAO acknowledges support from ARO grant W911NF2410401.

Data Availability Statement: Code for the cWGAN-based method presented here is available on GitHub at https://github.com/bshaddy/cWGAN_fire_arrival_time_inference_from_VIIRS_GOES_terrain, accessed on 20 Decemeber 2025. The training, testing, and validation datasets utilized in the present work are available at https://doi.org/10.5281/zenodo.17409951.

Acknowledgments: The authors acknowledge the Center for Advanced Research Computing (CARC) at the University of Southern California, USA for providing the computing resources that have contributed to the research results reported within this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Measurement Operator

The measurement operator M is approximated and applied to fire arrival times $\boldsymbol{\tau}^{(i)}$ to produce measurements $\bar{\boldsymbol{\tau}}^{(i)}$ in the following steps, where fire arrival times are measured in hours:

1. $\boldsymbol{\tau}^{(i)}$ is coarsened to a resolution of 375 m using a box kernel. This accounts for the coarser resolution of VIIRS data.
2. Four copies of the coarsened $\boldsymbol{\tau}^{(i)}$ are created and denoted as $\boldsymbol{\tau}_j^{(i)}, j = 1, ..., 4$. This accounts for the availability of new VIIRS measurements approximately two times per day, resulting in four distinct sets of measurements during the initial 48 h of a fire.
3. Components of $\boldsymbol{\tau}_j^{(i)}$ are eliminated or retained with a probability of $1/2$. This accounts for independent noise between individual VIIRS measurements.
4. Four measurement times $t_j, j = 1, ..., 4$ are sampled from a uniform probability distribution $U(2, max(\boldsymbol{\tau}^{(i)}) - 0.1)$ and sorted in ascending order. This enforces

randomness in the measurement collection time, increasing the variety in the resulting measurements.

5. For each measurement time t_j a time interval $(t_j - \delta, t_j)$ is generated, where δ is sampled from $U(6, 12)$, and negative values of $t_j - \delta$ are set to 0. This prescribes a random residence time for burning pixels, with the minimum selected to ensure a sufficient number of detections are captured [43].
6. For $j = 1, \cdots, 4$ fire arrival time values in $\boldsymbol{\tau}_j^{(i)}$ falling within the associated interval $(t_j - \delta, t_j)$ are set to t_j. Fire arrival time values outside of this interval are set to a background value. This allows pixels ignited prior to a VIIRS measurement which may still be burning (dictated by the prescribed residence time) to be captured.
7. Measurements $\boldsymbol{\tau}_j^{(i)}$ are combined into a single measurements by taking $\boldsymbol{\tau}^{(i)} = min_j(\boldsymbol{\tau}_j^{(i)})$ at each pixel. This ensures that the time when a burning pixel is first observed is assigned as the approximate fire arrival time.
8. An ignition time error δT is sampled from $U(0, 2)$ and subtracted from arrival time values in $\boldsymbol{\tau}^{(i)}$. The choice of this distribution was based on the difference between the reported ignition time and the first GOES measurement observed across 17 fires.
9. Two 3 km x 3 km patches are randomly eliminated from $\boldsymbol{\tau}^{(i)}$ and eliminated pixels are set to a background value. This accounts for persistent obstructions occurring across all VIIRS measurements.
10. All background pixel values are set to 48 h.
11. $\boldsymbol{\tau}^{(i)}$ is upsampled to the original resolution of 25 m and the resulting measurement is denoted by $\bar{\boldsymbol{\tau}}^{(i)}$.

Figure A1 additionally presents a flowchart summarizing the steps of the observation operator, with each transformation step categorized based on its overall effect on the fire arrival time map.

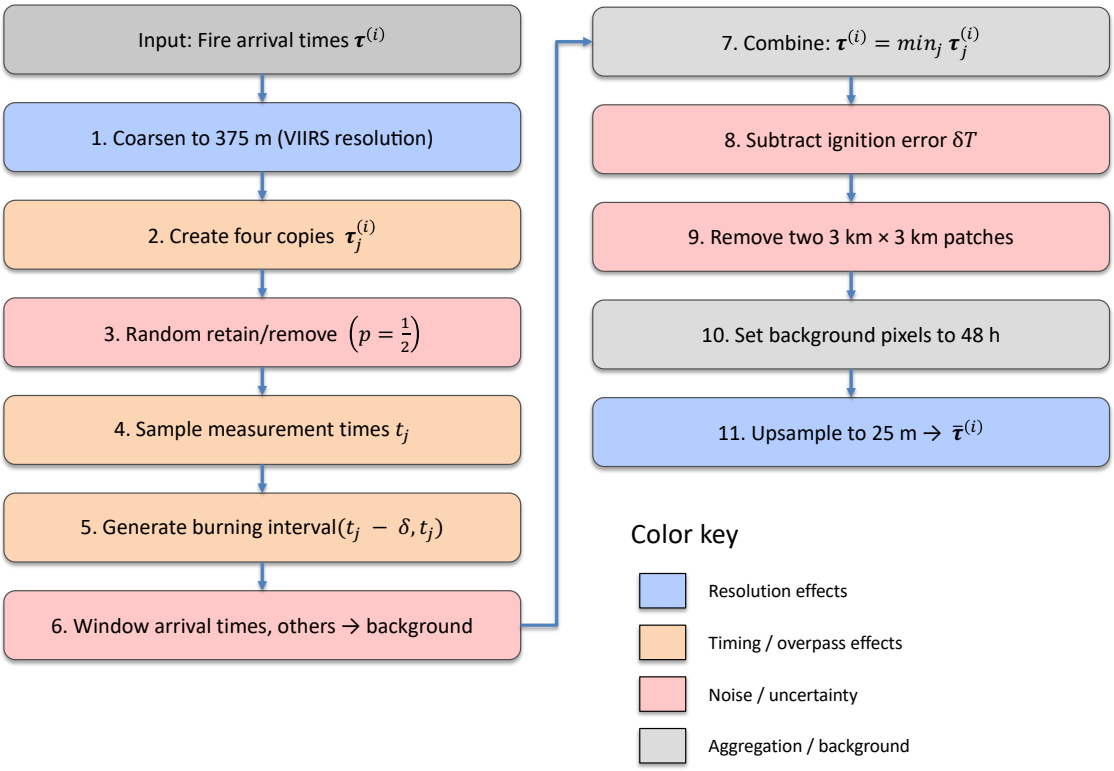


Figure A1. Flowchart highlighting the key transformations performed by the observation operator, with each step categorized as involving resolution, timing/overpass, noise/uncertainty, or aggregation/background.

Appendix B. cWGAN Architecture

The architectures used for the generator and critic subnetworks of the cWGAN are shown in Figure A2. The architecture follows that of Shaddy et al. [29], with the exception of an additional input channel to both networks in order to accommodate terrain as a conditional input. The generator uses a U-Net architecture, allowing for the extraction of features at multiple scales and the preservation of spatial details, and takes as input the measurement $\bar{\boldsymbol{\tau}}$ and terrain height $\boldsymbol{h}$, along with the latent vector $\boldsymbol{z}$. The generator is composed of dense blocks, downsampling blocks, and upsampling blocks, and uses conditional instance normalization (CIN) to inject latent information at multiple scales, overcoming issues of mode collapse [45,51]. The critic architecture is comprised of dense blocks, downsampling blocks, and fully connected layers, and takes as input tuples of $(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}, \boldsymbol{h})$ to produce a real-valued scalar.

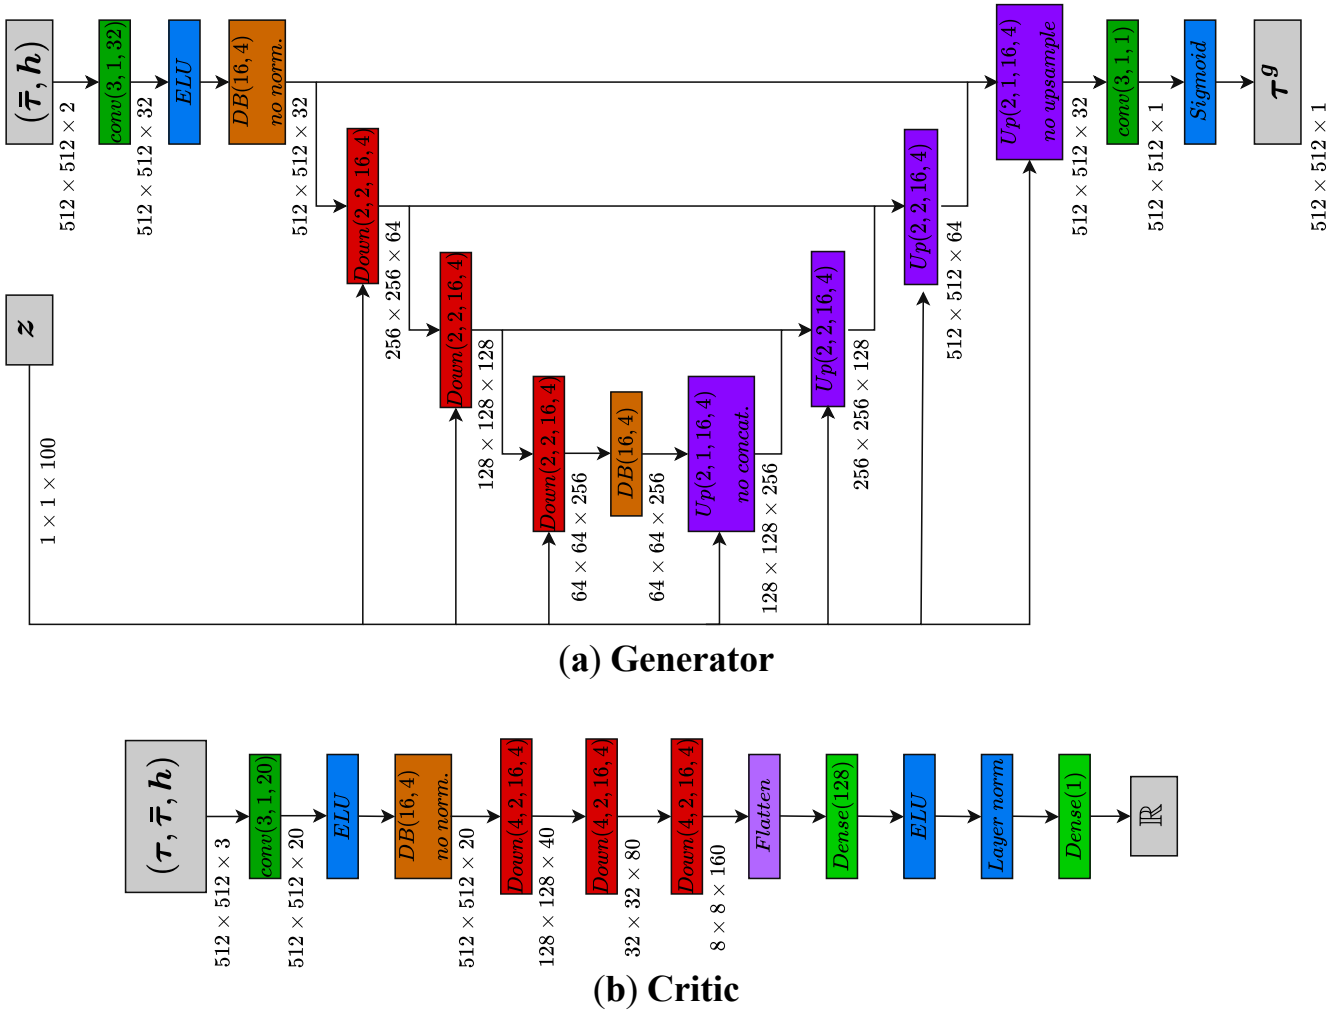


Figure A2. Network architecture for (**a**) generator and (**b**) critic of cWGAN.

Dense blocks follow the work of Huang et al. [52] and are denoted as $DB(k, n)$, where n is the number of sub-blocks and k is the number of output features from intermediate sub-blocks. Downsampling blocks consist of a convolutional layer, followed by Exponential Linear Unit (ELU) activation, average pooling, and a dense block, and are denoted as $Down(p, q, k, n)$, where p denotes the factor by which images are coarsened from average pooling, q denotes the factor by which channels are increased by convolutions, and k and n are the parameters of the dense block. Upsampling blocks are denoted as $Up(p, q, k, n)$ and consist of concatenation for skip connections, convolutions, ELU activation, interpolation,

and a dense block, where again p indicates the factor by which images are refined by interpolation, q is the factor by which channels are reduced, and k and n are the parameters for the dense block.

References

1. Dennison, P.E.; Brewer, S.C.; Arnold, J.D.; Moritz, M.A. Large wildfire trends in the western United States, 1984–2011. *Geophys. Res. Lett.* **2014**, *41*, 2928–2933. [CrossRef]
2. Juang, C.S.; Williams, A.P.; Abatzoglou, J.; Balch, J.; Hurteau, M.; Moritz, M. Rapid growth of large forest fires drives the exponential response of annual forest-fire area to aridity in the western United States. *Geophys. Res. Lett.* **2022**, *49*, e2021GL097131. [CrossRef] [PubMed]
3. Westerling, A.L.; Hidalgo, H.G.; Cayan, D.R.; Swetnam, T.W. Warming and earlier spring increase western US forest wildfire activity. *Science* **2006**, *313*, 940–943. [CrossRef] [PubMed]
4. Flannigan, M.; Cantin, A.S.; De Groot, W.J.; Wotton, M.; Newbery, A.; Gowman, L.M. Global wildland fire season severity in the 21st century. *For. Ecol. Manag.* **2013**, *294*, 54–61. [CrossRef]
5. Burke, M.; Driscoll, A.; Heft-Neal, S.; Xue, J.; Burney, J.; Wara, M. The changing risk and burden of wildfire in the United States. *Proc. Natl. Acad. Sci. USA* **2021**, *118*, e2011048118. [CrossRef]
6. Wang, D.; Guan, D.; Zhu, S.; Kinnon, M.M.; Geng, G.; Zhang, Q.; Zheng, H.; Lei, T.; Shao, S.; Gong, P.; et al. Economic footprint of California wildfires in 2018. *Nat. Sustain.* **2021**, *4*, 252–260. [CrossRef]
7. Danielle, M. AccuWeather Estimates More than $250 Billion in Damages and Economic Loss from La Wildfires. 2025. Available online: https://www.accuweather.com/en/weather-news/accuweather-estimates-more-than-250-billion-in-damages-and-economic-loss-from-la-wildfires/1733821 (accessed on 15 January 2025).
8. Lareau, N.P.; Clements, C.B. Environmental controls on pyrocumulus and pyrocumulonimbus initiation and development. *Atmos. Chem. Phys.* **2016**, *16*, 4005–4022. [CrossRef]
9. Lareau, N.P.; Clements, C.B. The mean and turbulent properties of a wildfire convective plume. *J. Appl. Meteor. Climatol.* **2017**, *56*, 2289–2299. [CrossRef]
10. Pastor, E.; Zárate, L.; Planas, E.; Arnaldos, J. Mathematical models and calculation systems for the study of wildland fire behaviour. *Prog. Energy Combust. Sci.* **2003**, *29*, 139–153. [CrossRef]
11. Ehrke, C.; Farguell, A.; Kochanski, A.K. Interactions Between a High-Intensity Wildfire and an Atmospheric Hydraulic Jump in the Case of the 2023 Lahaina Fire. *Atmosphere* **2024**, *15*, 1424. [CrossRef]
12. Benik, J.T.; Farguell, A.; Mirocha, J.D.; Clements, C.B.; Kochanski, A.K. Analysis of fire-induced circulations during the FireFlux2 experiment. *Fire* **2023**, *6*, 332. [CrossRef]
13. Bakhshaii, A.; Johnson, E.A. A review of a new generation of wildfire–atmosphere modeling. *Can. J. For. Res.* **2019**, *49*, 565–574. [CrossRef]
14. Mandel, J.; Beezley, J.; Kochanski, A. Coupled atmosphere-wildland fire modeling with WRF-Fire version 3.3. *Geosci. Model Dev. Discuss.* **2011**, *4*, 497–545.
15. Coen, J.L.; Schroeder, W. Use of spatially refined satellite remote sensing fire detection data to initialize and evaluate coupled weather-wildfire growth model simulations. *Geophys. Res. Lett.* **2013**, *40*, 5536–5541. [CrossRef]
16. Noonan-Wright, E.K.; Opperman, T.S.; Finney, M.A.; Zimmerman, G.T.; Seli, R.C.; Elenz, L.M.; Calkin, D.E.; Fiedler, J.R. Developing the US wildland fire decision support system. *J. Combust.* **2011**, *2011*, 168473. [CrossRef]
17. O'Connor, C.D.; Thompson, M.P.; Rodríguez y Silva, F. Getting ahead of the wildfire problem: Quantifying and mapping management challenges and opportunities. *Geosciences* **2016**, *6*, 35. [CrossRef]
18. Dunn, C.J.; Thompson, M.P.; Calkin, D.E. A framework for developing safe and effective large-fire response in a new fire management paradigm. *For. Ecol. Manag.* **2017**, *404*, 184–196. [CrossRef]
19. Thompson, M.P.; Wei, Y.; Calkin, D.E.; O'Connor, C.D.; Dunn, C.J.; Anderson, N.M.; Hogland, J.S. Risk management and analytics in wildfire response. *Curr. For. Rep.* **2019**, *5*, 226–239. [CrossRef]
20. Cao, Y.; Boruff, B.J.; McNeill, I.M. Is a picture worth a thousand words? Evaluating the effectiveness of maps for delivering wildfire warning information. *Int. J. Disaster Risk Reduct.* **2016**, *19*, 179–196. [CrossRef]
21. Schroeder, W.; Giglio, L.; Hall, J. *Collection 2 Visible Infrared Imaging Radiometer Suite (VIIRS) 375-m Active Fire Product User's Guide Version 1.0*; National Aeronautics and Space Administration (NASA): Washington, DC, USA, 2024.
22. Schmidt, C.; Hoffman, J.; Prins, E.; Lindstrom, S. *GOES-R Advanced Baseline Imager (ABI) Algorithm Theoretical Basis Document for Fire/Hot Spot Characterization, Version 2.0;* NOAA NESDIS Center for Satellite Applications and Research (STAR) : Silver Spring, MD, USA, 2010.
23. Kondratenko, V.Y.; Beezley, J.D.; Kochanski, A.K.; Mandel, J. Ignition from a fire perimeter in a WRF wildland fire model. *arXiv* **2011**, arXiv:1107.2675. [CrossRef]

24. Mandel, J.; Beezley, J.D.; Kochanski, A.K.; Kondratenko, V.Y.; Kim, M. Assimilation of perimeter data and coupling with fuel moisture in a wildland fire–atmosphere DDDAS. *Procedia Comput. Sci.* **2012**, *9*, 1100–1109. [CrossRef]
25. Mandel, J.; Kochanski, A.K.; Vejmelka, M.; Beezley, J.D. Data assimilation of satellite fire detection in coupled atmosphere-fire simulation by WRF-SFIRE. *arXiv* **2014**, arXiv:1410.6948. [CrossRef]
26. Kochanski, A.K.; Clough, K.; Farguell, A.; Mallia, D.V.; Mandel, J.; Hilburn, K. Analysis of methods for assimilating fire perimeters into a coupled fire-atmosphere model. *Front. For. Glob. Change* **2023**, *6*, 1203578. [CrossRef]
27. Mandel, J.; Amram, S.; Beezley, J.; Kelman, G.; Kochanski, A.; Kondratenko, V.; Lynn, B.; Regev, B.; Vejmelka, M. Recent advances and applications of WRF–SFIRE. *Nat. Hazards Earth Syst. Sci.* **2014**, *14*, 2829–2845. [CrossRef]
28. Farguell, A.; Mandel, J.; Haley, J.; Mallia, D.V.; Kochanski, A.; Hilburn, K. Machine Learning Estimation of Fire Arrival Time from Level-2 Active Fires Satellite Data. *Remote Sens.* **2021**, *13*, 2203. [CrossRef]
29. Shaddy, B.; Ray, D.; Farguell, A.; Calaza, V.; Mandel, J.; Haley, J.; Hilburn, K.; Mallia, D.V.; Kochanski, A.; Oberai, A. Generative Algorithms for Fusion of Physics-Based Wildfire Spread Models with Satellite Data for Initializing Wildfire Forecasts. *Artif. Intell. Earth Syst.* **2024**, 3, e230087. [CrossRef]
30. Veraverbeke, S.; Sedano, F.; Hook, S.J.; Randerson, J.T.; Jin, Y.; Rogers, B.M. Mapping the daily progression of large wildland fires using MODIS active fire data. *Int. J. Wildland Fire* **2014**, *23*, 655–667. [CrossRef]
31. Scaduto, E.; Chen, B.; Jin, Y. Satellite-based fire progression mapping: A comprehensive assessment for large fires in Northern California. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2020**, *13*, 5102–5114. [CrossRef]
32. Parks, S.A. Mapping day-of-burning with coarse-resolution satellite fire-detection data. *Int. J. Wildland Fire* **2014**, *23*, 215–223. [CrossRef]
33. Chen, Y.; Hantson, S.; Andela, N.; Coffield, S.R.; Graff, C.A.; Morton, D.C.; Ott, L.E.; Foufoula-Georgiou, E.; Smyth, P.; Goulden, M.L.; et al. California wildfire spread derived using VIIRS satellite observations and an object-based tracking system. *Sci. Data* **2022**, *9*, 249. [CrossRef]
34. Clough, K.; Farguell, A.; Mandel, J.; Hilburn, K.; Kochanski, A. Enhancing wildfire and smoke forecasting by integrating fire observations: A comparative analysis of methods for integrating infrared and satellite data into a coupled fire-atmosphere model. *J. Geophys. Res. Atmos.* **2025**, *130*, e2024JD042561. [CrossRef]
35. McCarthy, N.F.; Tohidi, A.; Aziz, Y.; Dennie, M.; Valero, M.M.; Hu, N. A deep learning approach to downscale geostationary satellite imagery for decision support in high impact wildfires. *Forests* **2021**, *12*, 294. [CrossRef]
36. Lai, C.H.; Song, Y.; Kim, D.; Mitsufuji, Y.; Ermon, S. The principles of diffusion models. *arXiv* **2025**, arXiv:2510.21890. [CrossRef]
37. Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; Zhu, J. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 5775–5787.
38. Zhou, Z.; Chen, D.; Wang, C.; Chen, C. Fast ode-based sampling for diffusion models in around 5 steps. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024 ; pp. 7777–7786.
39. Shih, A.; Belkhale, S.; Ermon, S.; Sadigh, D.; Anari, N. Parallel sampling of diffusion models. *Adv. Neural Inf. Process. Syst.* **2023**, *36*, 4263–4276.
40. Zheng, H.; Nie, W.; Vahdat, A.; Azizzadenesheli, K.; Anandkumar, A. Fast sampling of diffusion models via operator learning. In Proceedings of the International Conference on Machine Learning, Honolulu, HI, USA, 15–19 June 2023; pp. 42390–42402.
41. Rothermel, R.C. *A Mathematical Model for Predicting Fire Spread in Wildland Fuels*; Research Paper, INT-115. US Department of Agriculture, Intermountain Forest and Range Experiment Station: Ogden, UT, USA , 1972.
42. Mandel, J.; Vejmelka, M.; Kochanski, A.; Farguell, A.; Haley, J.; Mallia, D.; Hilburn, K. An interactive data-driven HPC system for forecasting weather, wildland fire, and smoke. In Proceedings of the 2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC), Denver, CO, USA, 17 November 2019; IEEE: Piscataway, NJ, USA , 2019; pp. 35–44.
43. Mandel, J.; Fournier, A.; Jenkins, M.A.; Kochanski, A.K.; Schranz, S.; Vejmelka, M. Assimilation of satellite active fires detection into a coupled weather-fire model. In *Proceedings of the 5th International Fire Behavior and Fuels Conference, Portland, OR, USA, 11–15 April 2016*; International Association of Wildland Fire: Missoula, MT, USA, 2016; pp. 17–22.
44. Ray, D.; Murgoitio-Esandi, J.; Dasgupta, A.; Oberai, A.A. Solution of physics-based inverse problems using conditional generative adversarial networks with full gradient penalty. *arXiv* **2023**, arXiv:2306.04895. [CrossRef]
45. Adler, J.; Öktem, O. Deep bayesian inversion. *arXiv* **2018**, arXiv:1811.05910. [CrossRef]
46. Ray, D.; Ramaswamy, H.; Patel, D.V.; Oberai, A.A. The efficacy and generalizability of conditional GANs for posterior inference in physics-based inverse problems. *arXiv* **2022**, arXiv:2202.07773. [CrossRef]
47. Johnson, J.; Finco, M.; Quayle, B. Wide area imager fire mapping evaluation mission. In *Proceedings of the RSAC-10097-RPT1*; U.S. Department of Agriculture, Forest Service, Remote Sensing Applications Center: Salt Lake City, UT, USA, 2014.
48. Greenfield, P.H.; Smith, W.; Chamberlain, D.C. Phoenix-the new forest service airborne infrared fire detection and mapping system. In Proceedings of the 2nd International Wildland Fire Ecology and Fire Management Congress and the 5th Symposium on Fire and Forest Meteorology, Orlando, FL, USA, 16–20 November 2003.

49. Song, Y.; Sohl-Dickstein, J.; Kingma, D.P.; Kumar, A.; Ermon, S.; Poole, B. Score-based generative modeling through stochastic differential equations. *arXiv* **2020**, arXiv:2011.13456.
50. Dasgupta, A.; da Cunha, A.M.; Fardisi, A.; Aminy, M.; Binder, B.; Shaddy, B.; Oberai, A.A. Unifying and extending diffusion models through pdes for solving inverse problems. *arXiv* **2025**, arXiv:2504.07437. [CrossRef]
51. Dumoulin, V.; Shlens, J.; Kudlur, M. A learned representation for artistic style. *arXiv* **2016**, arXiv:1610.07629. [CrossRef]
52. Huang, G.; Liu, Z.; Van Der Maaten, L.; Weinberger, K.Q. Densely connected convolutional networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 4700–4708. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.