GSMMC 2026 Group A: Fire Containment

Nuo Chen [1]†, Tajkera Khatun [2]‡, Terhi Nurminen [3], Solari Elee [4], Tomás Gillanders [5]§, Van Hoang Nguyen [6], and W. Hayden Taylor [7]

[1] *UCLA*
[2] *Indiana University Indianapolis*
[3] *University of Pennsylvania*
[4] *University of Delaware*
[5] *Mathematical Institute, University of Oxford, UK*
[6] *Texas Tech University*
[7] *University of Delaware*

(*Communicated to* MIIR *on June 2026*)

Study Group: Graduate Student Mathematical Modeling Camp 2026, University of Delaware.
Presenter: Dr. Manuchehr Aminian, Cal Poly Pomona

Summary

The modeling of wildfire spread and algorithmic approaches to containing the spread are considered. Various partial differential equation (PDE) models of how wildfires spread are considered, with varying levels of complexity. Numerical methods are used to solve these systems of PDEs to allow for the forecasting of how the wildfires spread. This information is then used along with identified fire containment algorithms to simulate how the fire's spread can be subdued. Some of these fire-containment approaches are qualitatively compared against one another. Additional complications to the modeling of the fire spread, such as stochastic fluctuations in the wind direction are also considered.

Contents

† Email: `nchen@math.ucla.edu`
‡ Email: `takhatun@iu.edu`
§ Email: `tomas.gillanders@maths.ox.ac.uk`

1 Problem Statement

In considering the problem of effective wildfire containment, the sponsor listed the following key points of interest that they wished to be elaborated on:

(1) Using a purely diffusive fire-spread toy model, how can one design and implement a fire containment model? How do different containment approaches compare with each other, and is there an approach to 'optimally' constrain the fire spread? Associated to this preliminary model, are there key/relevant nondimensional parameters that elucidate the behavior of the system?

(2) Can an enriched PDE model incorporating aspects of the physics of fire-spread be identified, incorporating at least one additional modeling characteristic? Using this model, can one implement a numerical PDE-solving scheme together with a containment algorithm that updates the containment path over time, as demonstrated in numerical experiments?

(3) Can methods for modeling and parameter estimation of fire spread in time with uncertainty be identified? This may entail applying filtering methods, stochastic

optimization for optimal containment, or variance reduction methods. Can toy models incorporating these ideas be identified?

(4) Using the models and containment methods considered by the group, can the containment procedure be compared with data pertaining to historical fires and their containment to determine the efficacy of the procedures identified? What metrics for the success of the containment algorithms can be identified?

2 Introduction

As anthropogenic climate change causes temperatures to rise, the threat of wildfires has grown significantly. Wildfires pose a significant danger to both natural environments and human lives and property; an uncontrolled forest wildfire can devastate massive swathes of precious natural land, disrupting the (already threatened) ecosystems there; Meanwhile, an urban wildfire can have a high death toll and leave many more financially ruined due to the loss of property. As fire weather grows more common, existing fire containment resources are increasingly strained. Accordingly, it is important to determine the most optimal use of these resources, which requires us to understand two key components: how do wildfires spread, and how can they be most effectively contained?

In this report, we investigate the existing literature on the subject, and propose a model for wildfire spread which includes the effects of fuel and wind, as well as a stochastic time-dependent wind formulation that captures temporal variations in wind speed and direction. We then investigate several approaches to fire containment, including new methods (or variants of existing methods) developed during the GSMMC workshop. The stochastic wind model enables comparisons of wildfire spread under constant and time-dependent wind conditions, providing insight into the influence of realistic atmospheric variability on fire propagation.

3 Literature Review

- Bressan [2007] goes into technical mathematical language to define a class of problems which we associate with "surrounding a fire"; mentions of countably many Lipschitz continuous curves with a notion of "fixed rate of effort" σ (our words not his).
- Zambon et al. [2018] describes an algorithm for containment based on defining an arbitrary polygonal boundary, restricting potential barriers to diagonals of the polygon, and selecting appropriate barriers to construct based on an integer programming formulation.
- Mandel et al. [2008] builds up a coupled PDE of temperature T and fuel F from which both numerical simulation as well as some mathematical analysis

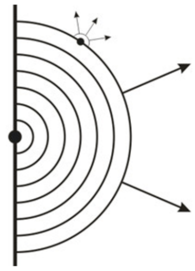

Figure 1. Schematic demonstrating Huygens' Principle, where the fire boundary propagates outwards from the source (black dot at center). Each point on the fire boundary at each instant also acts as a source of the fire.

4 Modeling Wildfire Spread

4.1 Huygens' Principle and Fire Boundary Propagation

A simple, "first approach", method to model fire boundary propagation is to consider the boundary as being an isotropically spreading wavefront from every point on the fire boundary (i.e. the fire emanates from every point on the current fire boundary and propagates in every direction not filled by fire at a fixed speed). This model of fire spread is analogous to that of Huygens' Principle from the theory of Geometric Optics (see Figure 1).

While this model ignores any variability in environmental conditions and any stochasticity or irregularities in the spatial spreading of the fire, its simplicity gives a natural mechanism for the fire to spread around barriers (such as natural environmental barriers or incomplete containment boundaries), and provides a useful toy model that can be used as a foundation with respect to which fire containment algorithms can be constructed. Having determined these fire containment algorithms, one could then aim to generalize the methods developed to work with more sophisticated models of wildfire spreading.

Such a model of fire spread is employed by Bressan [2007, 2013, 2026], Bressan et al. [2008], Bressan and Chiri [2022], Finney [1998, 2002]. Such an approach will also be taken in developing some of the algorithms of Section 6, before they are then adjusted for use with the more complicated models of fire spread described in the proceeding sections through use of numerical experiment.

4.2 Phenomenological PDE Models

We briefly outline a few common models which model fairly phenomenological aspects of fire spread.

4.2.1 *Classical Fisher-KPP Equation*

Phenomenological model which exhibits maxing out u at 1, supposing u represents "burning or burned":

$$u_t - \Delta u = ru(1-u) \tag{4.1}$$

Parameter Name	Description
$u(x, y, t)$	Fire density (normalized burn state, $0 \leq u \leq 1$)
t	Time (s)
f	Fuel, $0 \leq u \leq 1$)
$\mathbf{v} = (v_x, v_y)$	Wind (advection) velocity field (m/s)
D	Diffusion coefficient (fire spread intensity) (m^2/s)
r	Reaction rate coefficient (1/s)

Table 1. Parameter table for both models of Fisher-KPP (without Wind and Wind + Fuel Depletion + Containment)

4.2.2 *Fisher-KPP with Advection*

$$u_t + \mathbf{v} \cdot \nabla u - \Delta u = ru(1-u) \tag{4.2}$$

4.2.3 *Extended Fisher-KPP with Advection and Fuel*

Fire density equation $u(t, x)$

The model evolves two coupled fields $u(x, y, t)$ and $f(x, y, t)$ with diffusion, advection, and reaction terms to produce a realistic visual simulation of the fire spreading on a discrete grid:

$$u_t - D\Delta u + \mathbf{v} \cdot \nabla u = fu + ru(1-u), \tag{4.3}$$

$$f_t = -fu - ru(1-u). \tag{4.4}$$

- $D\,\Delta u$ is the diffusion term (embers carried by heat);
- $-\mathbf{v} \cdot \nabla u$ is advection by wind $\mathbf{v} = (v_x, v_y)$;
- $f\,u$ is ignition of fresh fuel by existing fire; The fuel equation decreases f wherever fire is present, ensuring eventual burnout inside fully consumed region.
- $r\,u(1-u)$ is a logistic (Fisher–KPP) reaction term modelling the self-sustaining spread of the combustion front.

See table 1 for a summary of parameters and variables.

4.3 Physics-Informed Reaction-Diffusion Model

In order to address Objective 2, as outlined by the sponsor, we adopt a simple reaction-diffusion model suggested by Mandel et al. [2008]. This physics-informed model incorporates the physical constraint of a finite fuel source (described by the fuel distribution, F) at each point in the forest domain that depletes as the fire burns, and an additional advective term, allowing for the incorporation of the effects of wind (described by a (possibly) time-dependent wind velocity, $\mathbf{u}$) in our fire-spread model.

In this model, we note that temperature, T, is used as a proxy for fire, where the distribution of the wildfire is determined by the regions where the temperature exceeds some threshold/ignition temperature, T_{ign}. Furthermore, the fuel distribution, F, is described as a proportion of the fuel remaining. Assuming that fuel is uniformly distributed in the domain of interest (i.e. we might assume that the forest area is densely packed with equally flammable and evenly distributed foliage), the fuel proportion, F, can take values between F $= 1$ (all of the fuel at a point is remaining) or F $= 0$ (none of the fuel at a point is remaining). We note that for the purposes of this report, we will always impose a uniform initial distribution of fuel of F $= 1$ across the domain of definition (i.e. the forest).

The equations describing the dynamic spreading of the wildfire are given by the following variation of the equations suggested by Mandel et al. [2008]:

$$\begin{aligned} \frac{\partial T}{\partial t} - D\Delta T + \mathbf{u}\cdot\nabla T &= -\lambda(T - T_a) + q\mathrm{F}\, r(T) \\ \frac{\partial \mathrm{F}}{\partial t} &= -\beta F\, r(T) \\ r(T) &= \frac{1}{1+\exp(-\sigma(T - T_{\mathrm{ign}}))} \end{aligned} \tag{4.5}$$

Note that (4.5) differs from the model of Mandel et al. [2008] in the choice of 'fire activation function', $r(T)$, chosen. The function $r(T)$ behaves like an activation for fire and fuel consumption in the neighborhood of a critical ignition temperature T_{ign}, whose sensitivity depends on the *ignition sensitivity parameter*, σ (see Figure 4.3). The original work Mandel et al. [2008] uses a qualitatively distinct version from this which behaves like an exponential decay for increasing T, which we did not spend much time exploring. However, we note that for the range of temperatures considered ($T > 0$), and also considering extremal cases of temperature (for example, the large absolute temperature limit, $T \to \infty$), we suspect that the discrepancy of choices of the activation function, $r(T)$, does not qualitatively change the behavior observed from the solutions of (4.5).

Parameters, variables, and their units for this model are defined in Table 2. The terms in (4.5) have the following physical interpretations:

- $-D\,\Delta T$: *Diffusion.* Allows for the qualitative spreading of the wildfire that we would expect.
- $\mathbf{u}\cdot\nabla T$: *Advection.* Here, we incorporate the wind, $\mathbf{u}$, as a mechanism for carrying the wildfire. Note that combining the time derivative and advection terms, we can define

$$\frac{\mathrm{D}T}{\mathrm{D}t} = \frac{\partial T}{\partial t} + \mathbf{u}\cdot\nabla T\,. \tag{4.6}$$

  This is familiar material time derivative of the temperature following the flow of the wind.
- $-\lambda(T - T_a)$: *Cooling.* The cooling rate of the fire is proportional to the temperature relative to the ambient temperature of the environment. This term can be interpreted as Newton's Law of Cooling [Davidzon, 2012].
- $qF\,r(T)$: *Heating.* The rate at which the temperature increases is proportional to the amount of fuel remaining and the rate at which the fuel is used (as determined by the activation function, r).
- $-\beta F\,r(T)$: *Fuel Usage.* The rate at which fuel is used is proportional to the amount

Parameter Name	Description
T	Absolute Temperature (K)
T_a	Ambient Temperature (K)
T_{ign}	Fire Ignition Temperature (K)
$\theta = T - T_a$	Temperature Relative to Ambient (K)
F	Fuel Proportion ($0 \leq \mathrm{F} \leq 1$)
D	Fire Diffusivity ($\mathrm{m^2\,s^{-1}}$)
$\mathbf{u} = (u_x, u_y)$	Wind Velocity ($\mathrm{m\,s^{-1}}$)
λ	Cooling Rate ($\mathrm{s^{-1}}$)
q	Fuel Heating Efficiency ($\mathrm{K\,s^{-1}}$)
β	Fuel Consumption Rate ($\mathrm{s^{-1}}$)
σ	Ignition Sensitivity Parameter ($\mathrm{K^{-1}}$)

Table 2. Variables and parameters of interest for the Mandel temperature and fuel model.

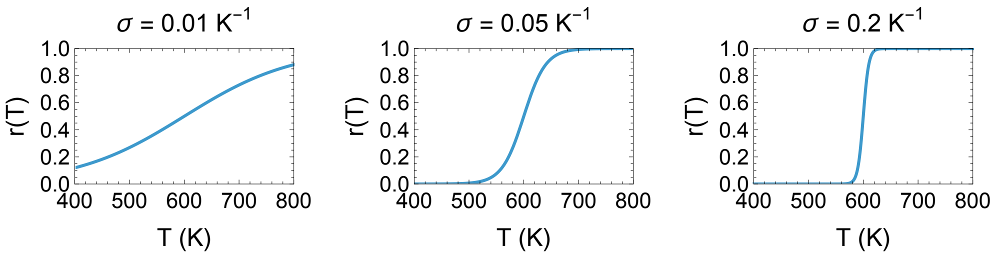

Figure 2. Fire activation function $r(T)$ for three values of the fire activation sensitivity parameter σ. The ignition temperature is fixed at $T_{\mathrm{ig}} = 600\,\mathrm{K}$.

of fuel remaining and the rate at which it is used (as determined by the activation function, r).

5 Numerical experiments

5.1 Fisher-KPP Equation

Two complementary experiments of the spatiotemporal spread of forest fire using the reaction-diffusion (and advection), where the fire intensity field $u(x, y, t) \in [0, 1]$ evolve on a 2D domain. The fire starts as a gaussian ignition in the center of the domain.

The first model has no depleting fuel only the reaction-diffusion-advection, so the logistic term acts as self-sustaining source wherever $u < 1$. There is no mechanism for fuel exhaustion. and so once the fire has been ignited the intensity monotonically increases, spreading out under the influence of reaction, diffusion and advection. This was also the case with no advection. Only much slower. This serves as controlled baseline for front propagation dynamics in the absence of resource limitation. The second model included a fuel density field $f(x, y, t) \in [0, 1]$ and fire intensity as coupled variables. The fuel actively depletes, as the fire intensity grows. Then we see a physically realistic burn scar.

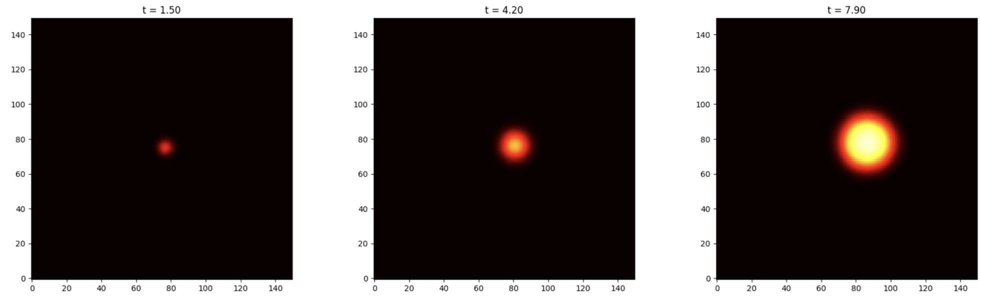

Figure 3. Numerical simulation of the solutions to the Fisher-KPP equation with an advective term (4.2).

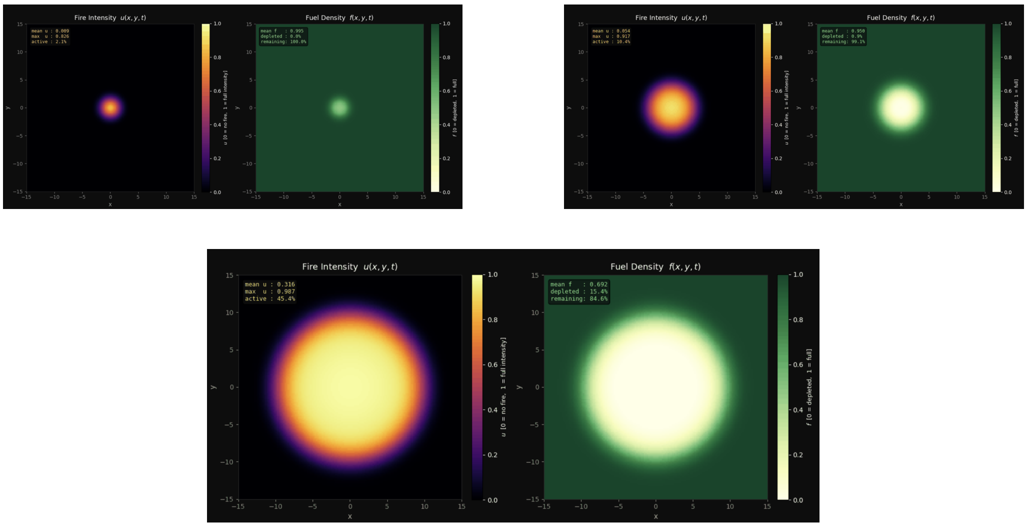

Figure 4. The images above show a close comparison of the fire generated by Fisher-KPP model without fuel and with fuel, respectively. In the Fisher-KPP model with fuel, when u = 0, no fire and f = 0, fully depleted fuel

5.2 Physics-Informed Reaction-Diffusion Model

We use this wildfire model as a common computational testbed for several containment strategies, including the tactical containment strategy, the Fried–Fried approach [Fried and Fried, 1996], the Zambon barrier-selection algorithm [Zambon et al., 2018], geometric containment constructions [Bressan and Wang, 2010], and the proposed SVD-based containment method. In each case, the fire evolution is governed by the same reaction-diffusion-advection fuel model, while the containment strategy modifies either the fuel field, the admissible barrier set, or the construction path used to stop the spread of the fire.

The computational domain is constructed from an observed wildfire perimeter obtained from CAL FIRE historical fire perimeter data [CalFire]. The perimeter is projected onto a local Cartesian coordinate system and discretized on a rectangular finite-difference grid. Grid points inside the observed fire polygon define the physical fire region, while grid points outside the polygon are masked in the visualization and diagnostic computations. The initial burned region is chosen near a selected ignition point inside the observed perimeter, and the initial temperature is prescribed as a localized hot region around this ignition set. Time-dependent wind data are obtained from Open-Meteo historical weather data [Zippenfenig, 2024] and converted into Cartesian velocity components ($\mathbf{u}(t) = (u_x(t), u_y(t))$), which are then used in the advection term of the PDE model. Writing $u = (T, F)^T$, the system has the semilinear form $u'(t) = Lu(t) + N(u(t))$, where

$$L\begin{pmatrix} T \\ F \end{pmatrix} = \begin{pmatrix} D\Delta T - \mathbf{u}\cdot\nabla T - \lambda T \\ 0 \end{pmatrix}, \qquad N(u) = \begin{pmatrix} \lambda T_a + qFr(T) \\ -\beta Fr(T) \end{pmatrix}.$$

If the wind field $\mathbf{u} = \mathbf{u}(t)$ is time-dependent, we put only the diffusion part into L and include the wind advection term in $N(t, u)$. For the spatial discretization, we use second-order centered finite differences on a uniform Cartesian grid, so for a two-dimensional grid function U, $\Delta_h U = D_{yy}U + UD_{xx}^T$. For time integration, we use a second-order exponential Runge–Kutta method with dimensional splitting, following the ETD2RKds idea in Caliari and Cassini [2024]. The stiff diffusion part is treated exponentially, while the reaction and wind terms are treated explicitly. The split matrix-function actions are computed direction by direction,

$$\Phi_k(W) = \varphi_k(\tau L_y)W\varphi_k(\tau L_x)^T, \qquad k = 1, 2.$$

This avoids forming the full two-dimensional Kronecker-sum matrix and gives an efficient implementation for the wildfire simulations.

6 Algorithms for Fire Containment & Results

In this section, we discuss a few different approaches to containing a wildfire. We begin by discussing existing analytical work and then implement and study various algorithms for containment.

6.1 Analytical Solutions to Isotropic Spreading

Given the large variability in environmental conditions and the stochastic and irregular features of wildfire spread, naturalistic simulations can typically not be studied using analytical methods. However, simplified setups often reveal features that are helpful in formalizing the research question and evaluating different approaches in the intractable scenario. Therefore, we include here a brief overview of a purely mathematical version of this problem first proposed by Bressan [2007], and discuss some known solutions.

6.1.1 *Problem setup*

Consider a fire that begins as a unit disc and propagates at unit speed according to Huygens' principle. The goal is to minimize the area the fire burns through by drawing a curve around the fire, starting from some point p at some constant velocity v. The curve is understood as an indestructible barrier that is impassable to the fire, with the only limitation being that any new segment of the curve has to be drawn on a region the fire

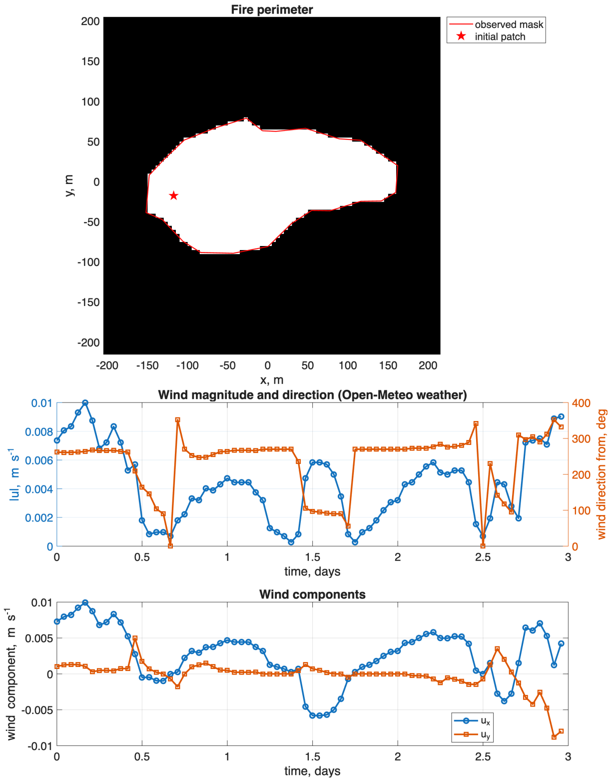

Figure 5. Observed wildfire perimeter and wind data used in the numerical simulation. Left: CAL FIRE observed fire perimeter projected onto a local Cartesian coordinate system, with the selected ignition point marked by a red star. Right: time-dependent Open-Meteo wind data converted into effective Cartesian velocity components. The top panel shows the wind magnitude $|\mathbf{u}|$ and wind direction, while the bottom panel shows the corresponding components $u_x(t)$ and $u_y(t)$ in m.s^{-1}.

has not yet reached. Notably, the curve is not required to be continuous or differentiable, and self-crossing is allowed, as is demonstrated in Figure 7.

In different variations of this problem, there may be one or multiple "firefighters" building barriers simultaneously, where for n firefighters their combined construction speed is v. Other variations include restricting the problem to the half plane with the fire originating as a half unit circle at the edge of the half plane [Bressan and Wang, 2009], existence of prebuilt barriers [Bressan and Chiri, 2022], and anisotropic spread [Bressan et al., 2008].

For a strategy to confine the fire within some region, a part of the curve has to eventually form either a closed loop, or, in the case of the half plane variant, be continuous from some point p_l on the edge of the half plane to another point p_r also on the edge of the half plane. Generally, as the goal is to minimize the burnt area, we look for solutions where the fire is contained in the bounded region delimited by the curve (and, where applicable, the edge of the half plane).

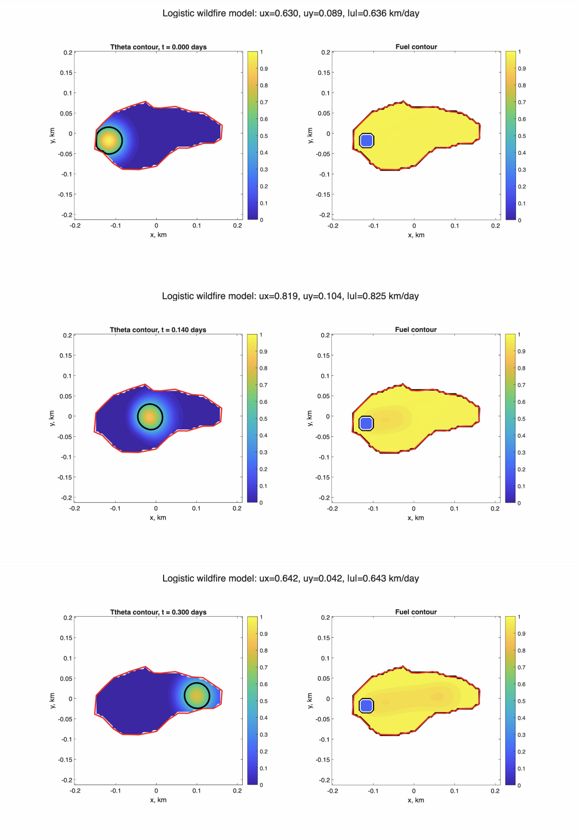

Figure 6. Snapshots of the reaction–diffusion–advection fuel model at selected times. For each time, the left panel shows the relative temperature field $\theta = T - T_a$, and the right panel shows the remaining fuel fraction F. The red curve denotes the observed CAL FIRE perimeter. The time-dependent wind velocity $\mathbf{u}(t) = (u_x(t), u_y(t))$ is shown above each snapshot and is used in the advection term of the PDE model.

6.1.2 *Existence of confinement strategies*

To begin our exploration of confinement strategies, we will first address the question of whether such strategies exist. One naive approach to the problem is to construct a perfectly circular barrier centered in the origin with one singular firefighter. For such a strategy, we are successful if at the time the curve is enclosed, denoted as t_e, the fire has spread at most to a distance equivalent to the radius of the circular barrier, r_c. This limiting case yields a system of equations

$$\begin{cases} \frac{2\pi r_c}{t_e} = v \\ \frac{r_c - 1}{t_e} = 1 \end{cases} \tag{6.1}$$

from which we can determine the smallest possible radius for the barrier curve:

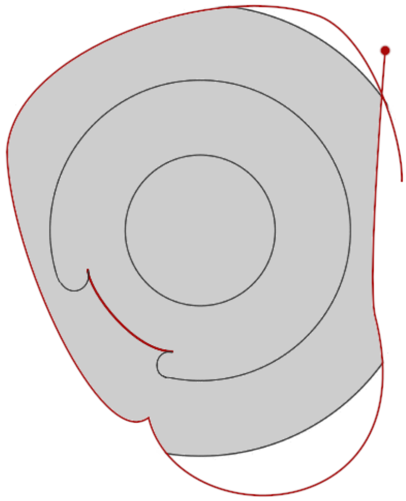

Figure 7. An illustration of an acceptable barrier curve successfully containing the fire. The barrier is not differentiable and features a self-intersection, as well as a disconnected component that affected the spread of the fire. The concentric black circles depict the spread of the fire over time, and the grey area represents the burned area at the time the berrier was finished.

$$r_c = \frac{v}{v - 2\pi} \tag{6.2}$$

Correspondingly, the minimum speed required to finish a curve with radius r_c is

$$v_{min} = \frac{2\pi r_c}{r_c - 1} \tag{6.3}$$

These results demonstrate that solutions to this problem exist, but more interestingly that there is a threshold speed under which a given geometry may not successfully contain the fire. Further, it should be noted that the problem is *scale agnostic*, with generalized expressions for r_c and v_{min} being

$$r_c = \frac{v r_0}{v - 2\pi} \tag{6.4}$$

$$v_{min} = \frac{2\pi r_c}{r_c - r_0} \tag{6.5}$$

where r_0 represents the radius of the fire at the start of barrier construction time t_0. In other words, if the starting condition is a scaled version of the original problem, the solutions can be scaled by the same factor without impacting existence or optimality. This applies for all strategies, not just the circular one showcased here.

This analysis suggests that the containment problem has two central components, namely, the speed required to construct a barrier and the optimal geometry for the barrier. At least the following questions arise:

(1) What is the critical speed below which it is not possible to construct a barrier that confines the fire?

(2) Do different construction speeds lead to different optimal geometries?

(3) Does the number of firefighters employed impact the critical speed, the optimal geometry, or their relationship?

The question of critical speed has been studied for a number of scenarios. For the general formulation of the problem, we know that it is not possible to contain the fire if the barrier is built at unit speed (which matches the propagation speed), and the current known upper bound for the critical speed is at $v > 2$ [Bressan, 2007], leaving a gap $1 < v \leq 2$ for which it is not known if solutions exist. When restricted to a half plane, it is simply sufficient that the construction speed is above unit speed [Bressan and Wang, 2009]. The strategies that achieve containment at the established critical speeds, as well as some other special cases, are discussed next.

6.1.3 *Optimality of confinement strategies*

Given that successful containment strategies exist, the question of optimality naturally follows. Bressan [2007] suggests that this should be viewed as an optimization problem over the (lost) value of the burned area and the cost of the barrier. In realistic scenarios, time to confinement could also be an independent factor, but analytical work has mainly focused on finding critical speeds for different setups and minimizing total area.

In conjunction with his initial formulation of the problem, Bressan [2007] also proved some necessary conditions for the shape of a singular area minimizing barrier: when constructed away from the edge of the fire, its shape must be an arc of a circle curving away from the fire, whereas when constructing at the fire boundary, it must be a logarithmic spiral. Further, the curve should be differentiable at all intermediate points. Regarding solutions with n firefighters and a continuous curve, he also found that the optimal case is axisymmetric and consequently the speed of each firefighter must be $\frac{v}{n}$ [Bressan et al., 2008]. Below we present a brief overview of best known solutions in various setups, which are also illustrated in Figure 8.

Confinement with a single firefighter, continuous curve

In the case of drawing a single connected curve, a strategy that traces the edge of the fire successfully contains it for a construction speed above $v_c = 2.6144...$ (see Bressan et al. [2008] for the exact value) and results in a tightening logarithmic spiral that eventually closes on itself (see Figure 8 a)). Among single spiral strategies, this critical speed is strict, but the logarithmic spiral is not unique in reaching it: some area-wise suboptimal deformed spirals achieve it as well [Bianchini and Zizza, 2025]. However, it is not known whether the logarithmic spiral is the area-minimizing shape.

Confinement with multiple firefighters, continuous curve

In a multifront version of the single spiral strategy, multiple logarithmic spirals are constructed at the edge of the fire simultaneously. Multifront strategies with $2n$ spirals are strictly better than a single spiral, with the 2 spiral solution shown in Figure 8 b) minimizing the burned area in that class of solutions [Bressan et al., 2008]. This strategy can be further improved by starting the barrier as an arc of a circle some distance from the fire and transitioning into the spirals when the edge of the fire reaches the firefighters [Bressan et al., 2008], which is shown in Figure 8 c). Figure 8 d) illustrates the difference between these two solutions for the same construction speed. The arc and spirals strategy is known to be optimal for barriers that consist of a single curve, achieving the lowest area of all strategies [Bressan et al., 2008, Bressan and Wang, 2012]. Further, the two spiral solution, with or without circle arc, is successful when $v > 2$, which is the lowest known speed at which containment is possible [Bressan, 2007].

Confinement in half plane

In a modified problem setup, the confinement region is restricted to a half plane with the fire originating from the origin as a half unit circle. Bressan and Wang [2009] found that in this geometry, confinement is possible if and only if $v > 1$, again using a a logarithmic spiral that traces the edge of the fire as shown in Figure 8 d). And similarly to the full plane solution, a better solution applies a circle arc first. It is not known, however, if this is an optimal solution on the half plane.

Non-traditional problem setups

The containment problem has also been investigated for anisotropic spread [Bressan et al., 2008] and for a case where the goal is to protect half of the full plane [Kim et al., 2021]. The former could be considered a toy model for spread under windy conditions, and the circle arc and spiral strategy of the isotropic case still appears optimal if started at the direction the fire spreads fastest to. The latter found that building delaying barriers (that do not contribute to the final containment) is found helpful in containing the fire. These barriers are illustrated in Figure 8 e).

Overall, while the original fire containment problem is still open, these restricted cases have revealed efficient ways to race the fire front, whether near it or at a distance. Currently, all but one of the best known solutions in various setups comprise of a single enclosure curve without internal, delaying arcs, though Kim et al. [2021]'s work suggests that clever placement of such barriers should be explored in different settings. A major shortcoming of the current analytical literature is the focus on critical speeds over examining strategies that are optimal in relation to some cost function, especially one where the cost of area and barrier are not uniform.

In relation to real world wildfires, Bressan's problem formulation is of course highly unrealistic. Therefore, it is most useful to view it as a perfect information and execution scenario under perfectly regular conditions. It provides information about conditions that make containment absolutely impossible, and on the other hand shows how barrier shape and construction speed are the key questions for successful containment, rather than for example the starting size of the fire. The model also makes the importance of fully containing the fire painfully clear: if there is a hole in the barrier, it can quickly make the rest of the wall obsolete as well. Finally, the model reveals how symmetricity impacts picking an optimal strategy, and, more relevant for actual wildfires, how crucial it is to choose your strategy based on the future edge of the fire, not its present location.

6.2 Free Burning Fire Boundary Containment Algorithm

6.2.1 *Modeling Approach to Wildfire Spread*

Fried and Fried [1996] suggested a simple algorithm for containing the spread of wildfires, based on the premise that the shape of a wildfire spreads in a self-similar manner about some given point. In an idealized setting where no external effects influence the spreading of the fire, Fried and Fried [1996] assumed that the fire spreads in a centrally symmetric manner (i.e. as a disc), as is described using Huygens' Principle in Section 4.1. However, when introducing effects such as a prevailing wind, Fried and Fried [1996], assume that the expanding shape of the fire distorts into some ellipse with eccentricity, ϵ, that expands about one of its foci.

Such an assumption about how the fire spreads may be justified qualitatively by observing how a simulated fire spreads under appropriate fire diffusivity and wind velocity conditions. However, it should be noted that the approach of Fried and Fried [1996] is not specific to the example of an elliptically spreading fire, and can be used for any self-similar spreading shape that surrounds a given fire.

6.2.2 *Description of Algorithm*

In the approach of Fried and Fried [1996], two teams of "fire containers" skirt around the boundary of the fire, each traveling in opposite directions (see Figures 9, 10 and 11). They describe two classes of two different approaches to contain a spreading fire; namely

- a *direct* attack strategy, where the "fire containment teams" remain on the fire boundary at all times (i.e. one foot in and one foot outside the fire) and cut off the fire as they move along the perimeter.
- a *parallel* attack strategy, where the "fire containment teams" add a buffer zone (of width, L) about the fire boundary and skirt the new perimeter in the same manner that they would, were it a direct attack.

For obvious reasons, it is not ideal for half a "fire container" to be in the fire when trying to contain it, therefore, we direct our attention to the *parallel* attack strategy.

Having chosen a parallel attack strategy to contain the fire, we are left with one remaining choice of which side of the fire to start the containment from; either at the tail of the fire (Figure 9) or at its head (Figure 10).

A high-level description of the Free Burning Fire Boundary Containment Algorithm, as described by Fried and Fried [1996], is presented below.[1] A more detailed version of the algorithm for a *direct tail attack* is presented in Appendix B.1. For a full description and a detailed discussion of the algorithm, refer to Fried and Fried [1996].

[1] Note that the pseudocode for this algorithm was written with the help of AI tools.

Algorithm 1 Free Burning Fire Boundary Containment Algorithm [Fried and Fried, 1996]

Require: Free-burning fire boundary $X(u,\epsilon), Y(u,\epsilon)$
Require: Boundary expansion function $h(t)$
Require: Upper and lower line-building rates $V_1(t), V_2(t)$
Require: Offset distance L
Require: Maximum simulation time $t_{\max}$
Require: Attack type $\tau \in \{\text{head}, \text{tail}\}$

Construct the Super Fire Boundary
For each point on the free-burning fire boundary, construct the outward normal direction.
Define the super free-burning fire boundary (super fbfb) as the locus of points located a distance L along the outward normal.

$$x = hX(u,\epsilon) + L\cos\psi, \qquad y = hY(u,\epsilon) + L\sin\psi,$$

where ψ denotes the outward normal angle.
Treat the super fbfb as the effective boundary to be contained.
Initialize Line Construction
if $\tau = \text{tail}$ **then**
    Select the tail of the super fbfb as the initial attack point.
else
    Select the head of the super fbfb as the initial attack point.
end if
Set the upper and lower line-building fronts Q_1 and Q_2 to coincide at the chosen attack point.
Set $t \leftarrow t_0$.
while $t \leq t_{\max}$ **do**
    Compute the local tangent and outward normal to the super fbfb.
    Advance the upper and lower line-building fronts according to the available line-building rates $V_1(t)$ and $V_2(t)$.
    Update the corresponding boundary parameters u_1 and u_2.
    if $\tau = \text{tail}$ **then**
        Advance u_1 counterclockwise from the tail.
        Advance u_2 clockwise from the tail.
    else
        Advance u_1 clockwise from the head.
        Advance u_2 counterclockwise from the head.
    end if
    Compute the separation distance

$$D = \|Q_1 - Q_2\|.$$

    if $D = 0$ **then**
        Record containment time t_c.

```
        return Contained Fire.
    end if
    if either line-building front cannot keep pace with expansion of the super fbfb then
        return Escaped Fire.
    end if
    Advance simulation time.
end while
return Escaped Fire.
```

6.2.3 *Comparing Tail and Head Containment Strategies*

Having the choice of whether to start at the head or tail of the wildfire, as it begins to spread, the natural question of which strategy does a better job of containing the fire (i.e. minimizing the area damaged by fire) arises.

A preliminary numerical experiment was carried out, where the (nondimensional) area contained by the algorithm (for both the tail and head strategies) and under identical conditions (same rate of fire spread and same speeds of "fire containers") was carried out. The results of this experiment are presented in Figures 9 and 10. Here, we observe that under the conditions considered, the head-on attack strategy was significantly more effective at containing the fire when compared to the tail-on strategy.

This suggests that in general, a head-on fire containment attack is more favorable/optimal to contain the fire. However, to justify this claim, further mathematical justification is required. We leave this to future work.

6.2.4 *Application to the Physics-Informed Reaction-Diffusion Model*

In order to apply the Free Burning Fire Boundary (FBFB) strategy to the more realistic model of fire spread (in this case, the Physics-Informed Reaction Diffusion Model introduced in Section 4.3), we take the following approach:

(1) Noting that the FBFB algorithm works for any shape of wildfire spread (see Figure 11), we perform a preprocessing fire-spread forecasting step, where the boundary of the fire is tracked over time.

(2) Having determined this fire boundary curve, and how it evolves in time, the FBFB algorithm (with appropriate alterations to account for the fact that the region may be translating as well as expanding) is applied until the fire is contained.

Here, we note that the success of the FBFB algorithm is heavily dependent on the speed of the "fire containers" relative to the rate at which the fire spreads. However, for sufficiently fast "fire containers", the fire will be contained.

Due to the constraint of time, this algorithm was not successfully implemented in its entirety in the simulations. However, a simplified version of this approach was implemented, where a bounding ellipse of the fire was identified in the preprocessing step and was then used as a basis for implementing the FBFB algorithm. The result of this simulation is presented in Figure 12.

6.3 The Geometric Firefighter's Algorithm

The geometric firefighter problem (GFP) described by Zambon Zambon et al. [2018] is an abstracted containment problem working under a slightly different paradigm. In the GFP, the problem is constrained to a polygonal domain (i.e the 'forest') and a spreading fire begins at a set point. A finite set B of potential containment barriers are predetermined, and the goal of the problem is to select an optimal subset of B that can be successfully constructed and which minimizes the area the fire burns. A barrier is considered successfully constructed if the entire barrier can be built before the fire touches any point of it. Klein et al. [2018] establishes that the GFP is NP-hard in general.

As the problem is confined to a domain, solutions to the GFP can make use of the boundaries of the region as a natural barrier (i.e representing the boundary of the forest, beyond which the fire has little to no fuel to continue spreading).

Generally, the set of barriers B is constrained in some way to make the problem computationally tractable. Usually the constraint requires that the barriers are straight line segments and the endpoints of the barriers are on the edges perimeter of the boundary. In the GFP algorithm described in Zambon et al. [2018], the barriers are further required to be diagonals of the polygon, though this is essentially equivalent to only requiring points on the boundary (as we can simply add trivial vertices to the polygon).

The Zambon algorithm treats this as an integer programming problem, optimizing the area saved over all possible barrier constructions subject to constraints. The algorithm accepts a polygon P, a set of barriers B, and a set of values δ such that δ_i is the first time the fire touches a barrier $b_i \in B$. The output is a subset of barriers B and the optimal order to build them in. The full details of the algorithm may be found in Zambon et al. [2018].

Our primary development to Zambon's work is integrating our numerical models into Zambon's GFP algorithm. Originally, the values δ are calculated assuming the fire is spreading uniformly from a point source. We instead determine δ by numerical simulation, which allows us to apply Zambon's algorithm to more complicated patterns of fire spread, such as those modeled in section 4.2.

We used Zambon's Disjoint GFP (DGFP) variant of the algorithm, which adds the further restriction that chosen barriers are not permitted to intersect. This significantly speeds up computation and is a simpler problem to handle (though Klein et al. [2018] proves that even the DGFP is NP-hard).

Currently, the implementation of the algorithm has an error where redundant barriers are not accounted for properly. This has been fixed in newer versions of the code, but all results displayed in the report are still using the original flawed code, resulting in some clear inefficiencies. An implementation of the algorithm on a simple test polygon (and the barriers selected as output) are depicted in figure 13, and the output of Zambon integrated with our numerical simulation is shown in figure 14 (which unfortunately shows the error, as it has two redundant barriers). This error is fixed in the newest versions of the Python implementation.

The Zambon algorithm is something of an outlier in containment methods discussed in this report; where the others rely upon constructing a continuous curve to enclose the fire and do not explicitly adjust for the boundaries of the forest, Zambon's algorithm acts by essentially blocking off burnable areas and exploiting the boundary of the forest as a natural containment barrier. Notably, the Zambon procedure does not have hard requirements on the speed of containment (relative to speed of fire); other algorithms may fail if the speed of containment is too slow to keep pace with the expanding flames, but this GFP abstraction will always attempt to salvage what it can, suggesting this approach may be a useful fallback technique for fast-moving wildfires that cannot be fully enveloped.

6.4 "Ice Cream Cone" Containment Algorithm

Here we outline the algorithm described in Bressan and Wang [2010] and discuss some interpretation of this.

Algorithm 2 "Ice Cream Cone" Algorithm

Require: Drift parameter $\kappa > 1$, construction speed σ, cost weight λ, initial vertex count n

1: **Define the minimum-time function**

$$T(x) = \frac{(\kappa x_1 + 1) - \sqrt{(\kappa x_1 + 1)^2 - (\kappa^2 - 1)\left(\frac{x_2^2}{1+x_2^2} - 1\right)}}{\kappa^2 - 1}.$$

2: **Initialize geometry**

3: $\theta \leftarrow \dfrac{2\pi}{n}$

4: Define radial distances

$$\boldsymbol{\rho} = (\rho_1, \rho_2, \dots, \rho_n).$$

5: Set initial guess by placing vertices on a vertical segment at

$$x_1 = \frac{16\sqrt{2}}{3} - 3 \approx 4.54.$$

6: **while** $n \leq N_{\max}$ **do**

7:     **repeat**

8:         Compute segment lengths $\|S_k\|$ for all polygonal edges.

9:         Compute fire-arrival times $T(S_k)$.

10:         Determine permutation

$$\alpha = (i_1, \dots, i_n)$$

such that

$$T(S_{i_1}) \leq T(S_{i_2}) \leq \dots \leq T(S_{i_n}).$$

11:         Compute $A \leftarrow frac12 \sin(\theta) \sum_{k=1}^{n} \rho_k \rho_{k-1}$

12:         Compute total barrier length $L \leftarrow \sum_{k=1}^{n} \|S_k\|$

13:         $f \leftarrow A + \lambda L$

Verify admissibility:

$$\text{BuiltLength}(t) \leq \sigma t \qquad \forall\, t.$$

Update $\boldsymbol{\rho}$ along the fixed rays using a line-search method to reduce f while remaining admissible.

until convergence

if a local minimum is reached **then**

Refine mesh:

$$n \leftarrow 2n.$$

end if

end while

return Optimized polygonal confinement barrier.

In the initial phase, the fire begins and the logistic term $ru(1-u)$ dominates the diffusion term rapidly increasing the intensity of the fire leading to saturation point ($u \mapsto 1$). There is no wall construction occurs in this phase.
Then the wall construction phase starts as the first barrier segment is laid at a distance close (but feasible) enough to the origin, and segments are added gradually forming a circular enclosure. From a distance, the enclosure looks like a full ring. It is important to note that this "circular" shape arises as a result of the shape of the fire. With stronger advection in one direction, the barrier built will have a shape that earns the title of the algorithm.

6.5 SVD Containment Algorithm

Up to this point, the containment strategies considered made use of either a fixed containment geometry or solved a global optimization problem that makes use of a preprocessing forecasting step. In this section, we describe a completely local geometric strategy, for which a "fire container" keeps track of the orientation of the fire boundary closest to them, and advances in their efforts to contain the fire subject to this local orientation of the boundary and a predetermined "safe distance" at which they should stay away from the fire.[2] The benefit of such a strategy is that the "fire container" can update their movement in real time and adapt to the changing fire boundary as it evolves, without ever considering the global behavior of the wildfire.

In order to make estimations about the local geometry of the boundary (relative to the "fire container") a Singular Value Decomposition (SVD) is used in the following way:

(1) Consider a collection of the closest (discrete) perimeter points to the "fire container" at a given instant.

(2) Compute the centroid of these points and subtract it from each of the points.

(3) Compute the SVD of the resulting set of points (arranged into a $2 \times n$ matrix, where n is the total number of points selected).

[2] This strategy was suggested by our mentor, Dr. Manuchehr Aminian.

(4) The dominant (right) singular vector that results is an estimate for the tangent to the boundary closest to the "fire container", while the second most dominant singular vector corresponds to the normal to the boundary.

Having determined the estimated directions of the tangent and normal to the fire boundary, the "fire container" can then determine the direction in which the containment barrier is advanced/built by combining tangential motion to the fire boundary with a correctional term that maintains the desired "safe distance" to the fire. Schematics for the information obtained from the SVD and how the fire containment perimeter progresses in time due to this algorithm are presented in Figure 16. (Note that in the example shown, the fire perimeter grows linearly in time, with a pre-described, yet otherwise arbitrary, shape.)

A high-level description of the SVD Containment Algorithm is given in the following.[3] For a more detailed description of the algorithm, see Appendix B.2.

Algorithm 3 SVD Containment Algorithm

Require: Forecasted fire perimeter $\{X(t_i)\}_{i=0}^{N}$
Require: Initial containment point r_0
Require: Step size γ
Require: Number of nearby perimeter points k
Require: Desired stand-off distance d^*
Require: Weighting parameter σ
for $i = 1$ to N **do**
    Select the k nearest forecasted perimeter points near r_{i-1}
    Compute the local centroid x_{bdry}
    Form the centered point cloud and compute its SVD
    Set $v_{\|}$ and $v_{\perp}$ to the local tangent and normal directions
    Compute the distance d from r_{i-1} to the local boundary
    Blend the tangent and normal directions to obtain a containment direction w
    Update the containment point:

$$r_i \leftarrow r_{i-1} + \gamma w$$

end for
return containment path $\{r_i\}_{i=0}^{N}$

7 Comparison of Containment Algorithms

7.1 Model Metrics

Due to time constraints we were unable to do a proper comparison of various containment strategies; however, we include here a brief discussion on how such a comparison would be done. In general, the 'effectiveness' of a containment strategy may be graded on certain

[3] This pseudocode was written with the help of AI tools.

metrics. The most appropriate metrics depend on the specific local conditions of the wildfire, but common metrics include:

- Area burned/saved
- Value of area burned
- Time until total containment
- Monetary/resource cost of containment

Our preferred metric is area burned; a containment algorithm is better if we restrict the wildfire to burning a smaller area. Although in specific scenarios we may prefer other metrics (for example, when firefighting resources are heavily strained we may prefer algorithms with a low total cost of containment), area serves as a good general metric for comparison that is easy to compute. In section 7.2 we include the relevant mathematics for effectively comparing containment algorithms by area.

7.2 Comparison of Containment Area via Green's Theorem

Let a region of containment R be parameterized via $r(t) = \langle x(t), y(t) \rangle$ for $t \in [a, b]$. Then, by Green's theorem, we can compute the area of this region via the integral:

$$A = \frac{1}{2} \int_a^b x(t)y'(t) - x(t)y'(t)dt.$$

Since our regions are described via a large finite collection of points, we can approximate the areas very closely via the trapezoidal rule:

$$A \approx \Delta t \left(\frac{f(t_1) + f(t_N)}{2} + \sum_{k=2}^{N-1} f(t_k) \right),$$

where t_k is a finite sequence of points in time where $x(t_k), y(t_k)$ are the coordinates of each given point on the boundary of the region and $f(t) = x(t)y'(t) - x(t)y'(t)$. So, by selecting a certain contour of the fuel plots in the aftermath of the wildfire under each containment algorithm, we can measure how each of the containment algorithms compare in reducing the burned area.

8 Stochastic Wind Modeling

8.1 Time-Dependent Wind Modeling

Wind is one of the most important environmental factors affecting wildfire spread. Strong winds can accelerate the propagation of a fire, while changes in wind direction can alter the shape and orientation of the fire front. Unlike the constant-wind assumption used in the previous section, the present model allows the wind field to evolve in time, providing a more realistic representation of atmospheric variability during wildfire spread. To capture these effects, the wildfire model incorporates a stochastic time-dependent wind field whose speed and direction evolve throughout the simulation [Pastor, 2003]. The stochastic wind model extends the wind velocity field used in the physics-informed reaction–diffusion model (Eq. (4.5)), while the governing wildfire equations remain unchanged.

The wind velocity vector is represented as

$$\mathbf{u}(t) = (u_x(t), u_y(t)) \tag{8.1}$$

where

$$u_x(t) = U(t)\cos(\phi(t))$$
$$u_y(t) = U(t)\sin(\phi(t))$$

Here, $U(t)$ denotes the wind speed and $\phi(t)$ denotes the wind direction angle. These velocity components enter the wildfire temperature equation through the advection term acting on the relative temperature $\theta = T - T_a$,

$$-u_x \frac{\partial \theta}{\partial x} - u_y \frac{\partial \theta}{\partial y} \tag{8.2}$$

which transport heat in the direction of the prevailing wind. Consequently, larger wind speeds lead to faster heat transport and potentially faster fire spread.

8.1.1 *Wind Speed Dynamics*

The temporal evolution of wind speed is modeled using an Ornstein–Uhlenbeck (OU) process. The OU process is a mean-reverting stochastic process that combines random fluctuations with a tendency to return toward a prescribed average wind speed [Obukhov et al., 2021]. The governing stochastic differential equation is

$$dU = \gamma(U_{\text{mean}} - U)dt + \sigma_U dW_t \tag{8.3}$$

where:

- U_{mean} is the long-term average wind speed,
- γ is the mean-reversion coefficient,
- σ_U controls the magnitude of random fluctuations,
- W_t is a standard Wiener process.

The first term forces the wind speed toward the average value U_{mean}, while the second term introduces stochastic variability representing atmospheric turbulence and gusts.

For numerical implementation, the OU process is discretized as

$$U^{n+1} = U^n + \gamma(U_{\text{mean}} - U^n)\Delta t + \sigma_U \sqrt{\Delta t}\, \xi^n \tag{8.4}$$

where ξ^n is a normally distributed random variable with mean zero and variance one.

The parameter values used in the simulations are:

- Mean wind speed: $U_{\text{mean}} = 0.20$ km/hr
- Mean-reversion coefficient: $\gamma = 2.0$ day^{-1}
- Fluctuation strength: $\sigma_U = 0.06$

Wind conditions are updated every six hours during the simulation.

8.1.2 *Wind Direction Dynamics*

In addition to speed fluctuations, the wind direction is allowed to vary over time. The prevailing wind direction is assumed to be northward, corresponding to

$$\phi_{\text{mean}} = \frac{\pi}{2} \tag{8.5}$$

To generate smooth directional changes, random angular perturbations are sampled from a von Mises distribution. The von Mises distribution is often referred to as the circular analogue of the normal distribution and is widely used for directional data [Carta et al., 2008, Al Yammahi et al., 2021].

The direction update is given by

$$\phi^{n+1} = \phi^n + \alpha\eta^n \tag{8.6}$$

where

$$\eta^n \sim \text{von Mises}(0, \kappa) \tag{8.7}$$

The parameter α controls the magnitude of directional changes, while κ determines how strongly the direction remains concentrated around the prevailing direction.

The values used in this study are:

- $\alpha = 0.05$
- $\kappa = 20$

A relatively large value of κ results in small directional deviations and therefore smooth changes in wind direction over time.

8.1.3 *Wind Vector Reconstruction*

After updating the wind speed and direction, the wind components are reconstructed as

$$u_x = U\cos(\phi)$$
$$u_y = U\sin(\phi)$$

These components are then incorporated into the wildfire model and influence heat transport, ignition patterns, and overall fire-front evolution.

8.1.4 *Physical Interpretation*

The combined Ornstein–Uhlenbeck and von Mises framework provides a realistic representation of atmospheric wind forcing. The OU process captures intermittent gusts while preventing unrealistic long-term growth of wind speed. The von Mises process generates smooth directional variability while maintaining a preferred prevailing direction. Together, these mechanisms produce a dynamically evolving wind field that more closely resembles real atmospheric conditions and allows the wildfire model to simulate realistic fire-spread behavior under changing weather conditions.

Since wildfire containment depends strongly on predicting the evolving fire perimeter, incorporating stochastic wind forcing provides a more realistic framework for assessing containment performance under uncertain environmental conditions.

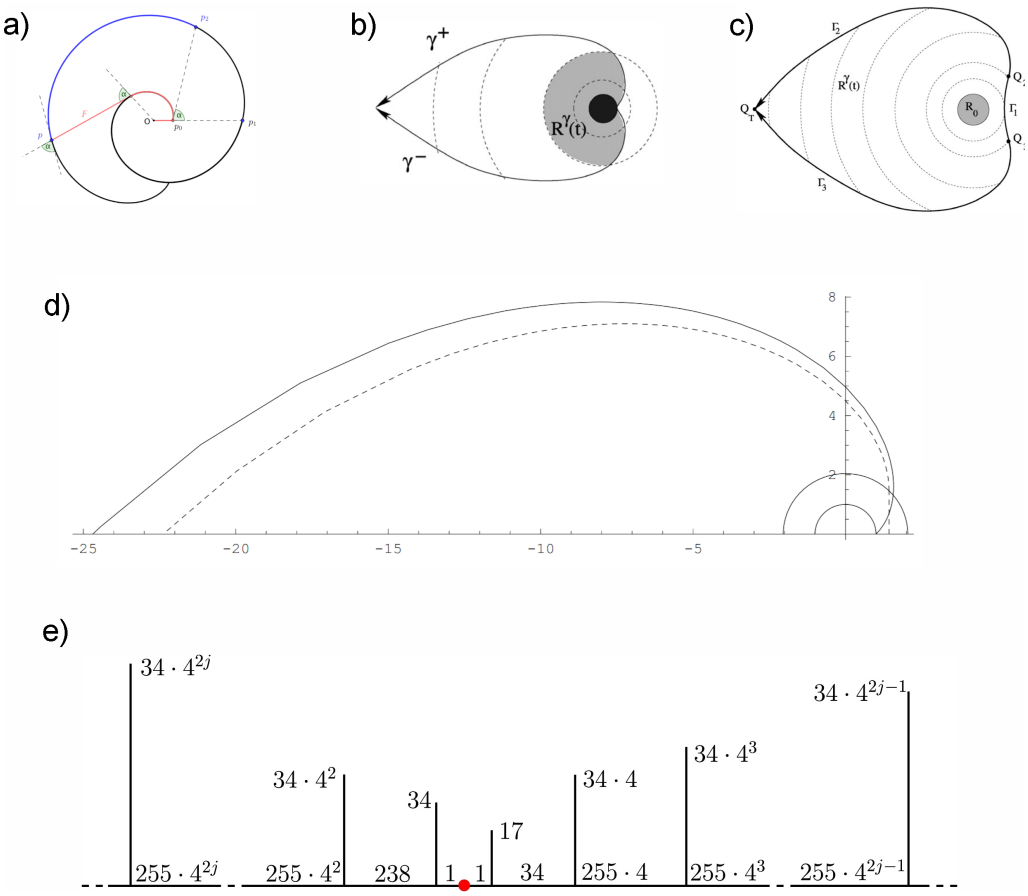

Figure 8. Compilation of best known strategies for fire containment with different problem setups. **a)** A tightening logarithmic spiral that traces the fire front. Best known solution for a single firefighter. Sketch from Klein et al. [2019]. **b)** Two logarithmic spirals tracing the fire front. The best multi firefighter strategy only utilizing spirals. Sketch from Bressan [2007] **c)** A circle arc combined with two logarithmic spirals. The optimal solution for the problem when applying a single curve. Sketch from Bressan [2007]. **d)** A circle arc and logarithmic spiral is the optimal strategy in half plane. This figure also demonstrates how using a circle arc as a starting shape (dashed line) results in a confinement with smaller area than the pure spiral arcs strategy (solid line). The original sketch in Bressan et al. [2008] is used to illustrate this area gain in the full plane scenario, but the same is true for half plane. **e)** An infinite straight line with perpendicular delaying segments. The best known strategy for preventing fire from reaching on half of a full plane, and the only known strategy where delaying segments decrease the critical speed. Sketch from Kim et al. [2021].

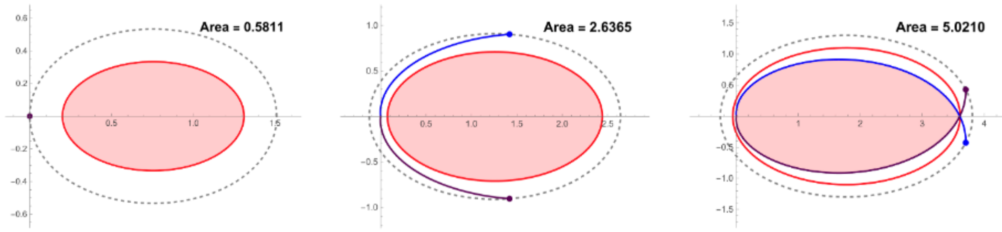

Figure 9. Parallel tail containment strategy with each of the "fire containers" traveling at equal speeds about the elliptical fire region, in opposite directions.

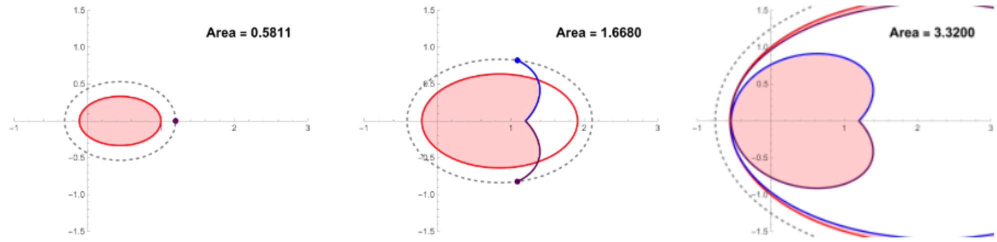

Figure 10. Parallel head containment strategy with each of the "fire containers" traveling at equal speeds about the elliptical fire region, in opposite directions.

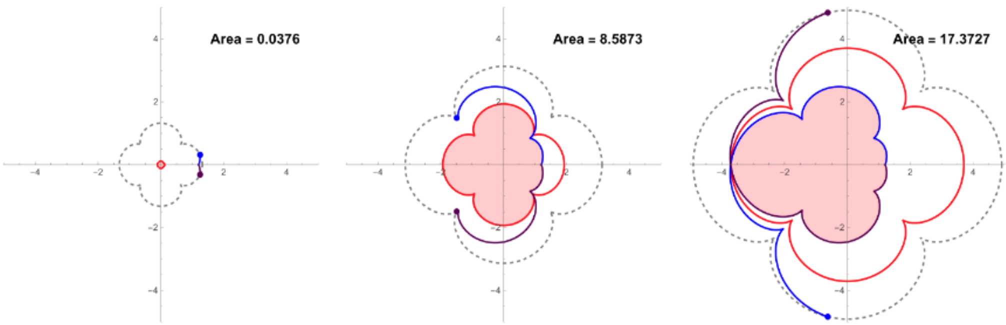

Figure 11. Free Burning Fire Boundary containment for an arbitrary self-similar fire spreading region.

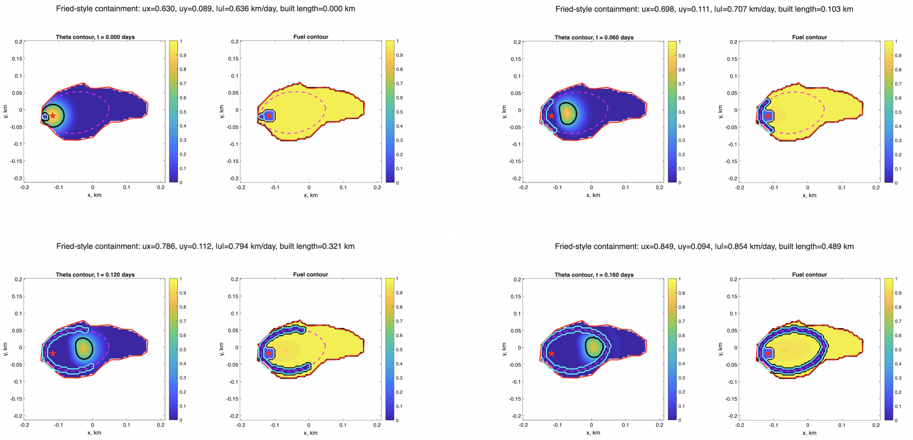

Figure 12. Free Burning Fire Boundary containment applied to the Physics-Informed Reaction-Diffusion model. Each snapshot shows the relative temperature field, $\theta = T - T_a$, and the remaining fuel fraction, F, at a selected time. The red curve denotes the observed fire perimeter and the cyan curve represents the constructed containment line. The magenta dashed curve shows the target containment path used by the FBFB strategy. The title reports the instantaneous wind velocity and the accumulated constructed containment length.

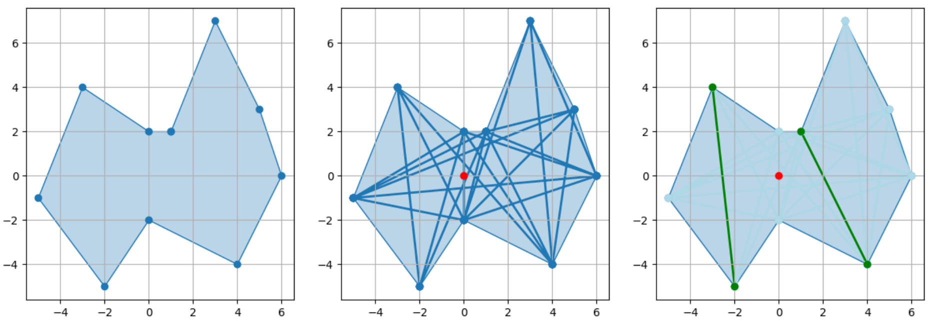

Figure 13. Depiction of Zambon's DGFP Algorithm. The first image shows a example polygon, the second highlights all possible barriers the algorithm chooses from, and the third shows the barriers selected by Zambon's DGFP algorithm.

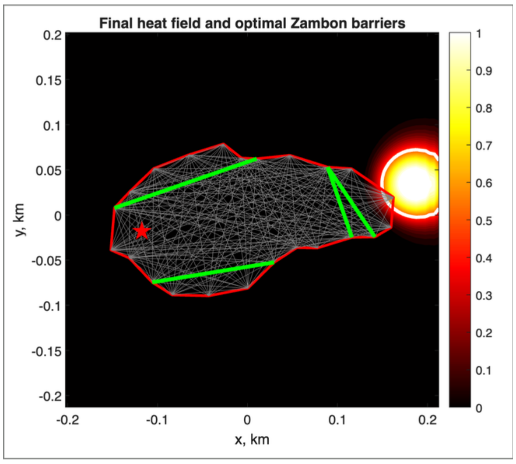

Figure 14. Output of Zambon DGFP Algorithm integrated with numerical simulation; the green lines represent the barriers selected by the algorithm. Due to a bug in the implementation, the algorithm selected the rightmost barriers incorrectly.

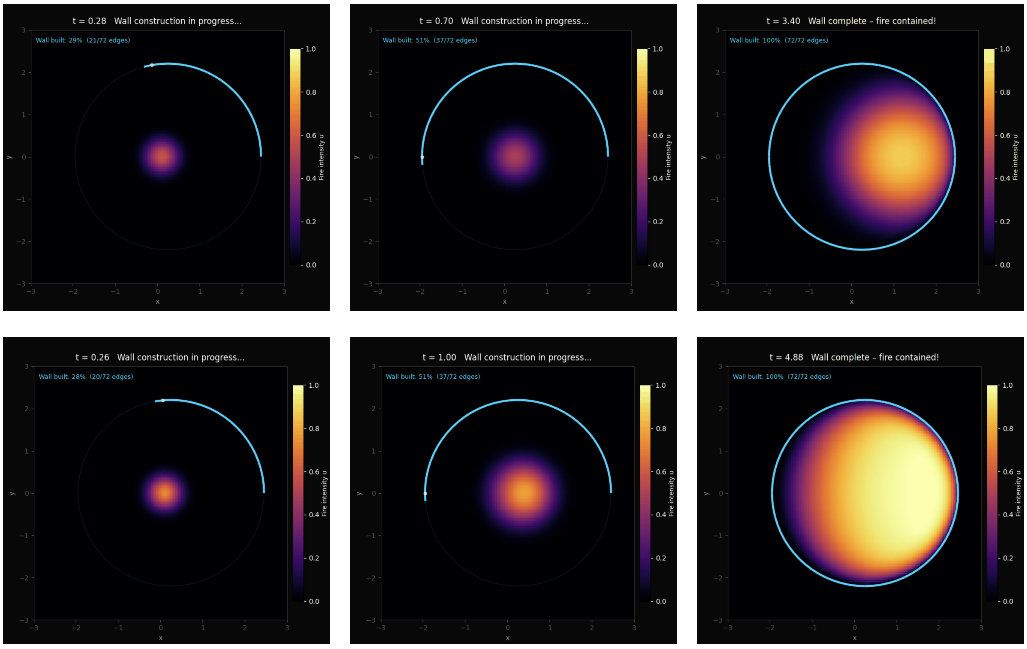

Figure 15. The fire containment model, ice cream cone, is coupled with the Fisher-KPP model, with and without fuel, respectively.

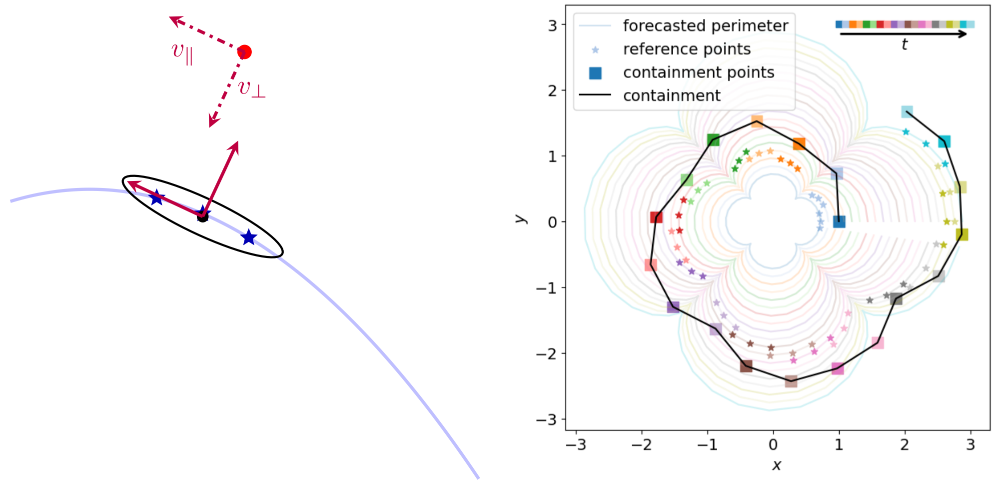

Figure 16. (Left) Illustration of local geometry informing a one-step choice of containment direction, computed using a SDV. (Right) SVD Containment Algorithm application with parameters $\gamma = 2\pi + 1$ (containment speed), $k = 3$ nearest perimeter points, target distance $d^* = 0.4$, and weighting falloff parameter $\sigma = \gamma$. See Appendix B.2 for more information about the parameters used.

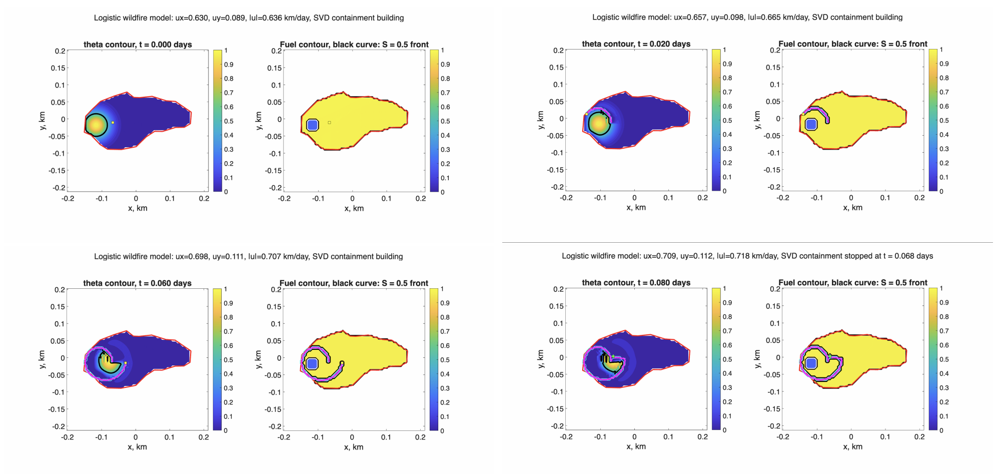

Figure 17. Time evolution of the logistic wildfire model with the proposed SVD-based containment construction. For each displayed time, the left panel shows the fire-intensity variable θ, while the right panel shows the remaining fuel F. The containment process stops the spread at approximately $t = 0.068$ days under the given time-dependent wind field

8.2 Simulation Results

The following figures compare the wildfire evolution under constant and stochastic time-dependent wind conditions. The constant-wind case serves as a baseline, while the stochastic case incorporates Ornstein–Uhlenbeck wind-speed fluctuations and smooth von Mises wind-direction changes. The wildfire model, numerical scheme, initial conditions, and containment configuration remain unchanged from the constant-wind simulations. Only the wind field is replaced by the stochastic model described above, allowing the influence of time-dependent wind variability on wildfire spread and containment performance to be isolated and directly compared.

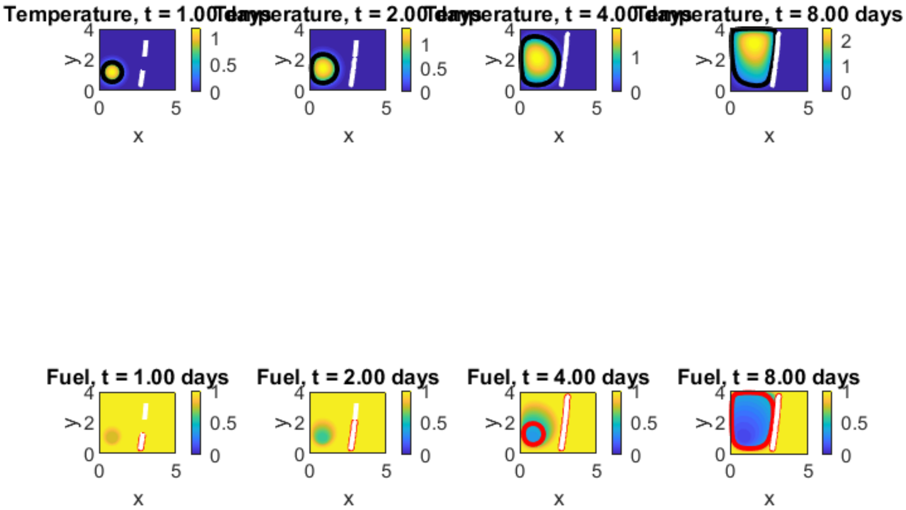

Figure 18. Temporal evolution of wildfire temperature and fuel distributions under constant wind conditions. The top row shows snapshots of the temperature field, while the bottom row shows the corresponding remaining fuel field at $t = 1$, 2, 4, and 8 days. The black contour denotes the ignition-temperature contour defining the fire front, and the red contour outlines the burned region. The white dashed line represents the containment line. A constant wind speed of $U = 0.20$ and a fixed wind direction of $\phi = 90^\circ$ are maintained throughout the simulation.

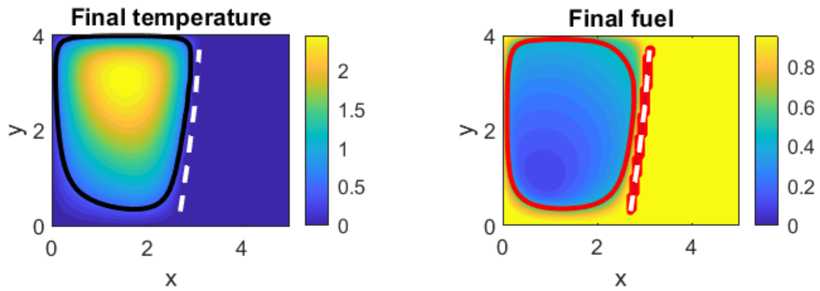

Figure 19. Final wildfire temperature and fuel distributions for the constant-wind simulation. The left panel shows the final temperature field, while the right panel shows the remaining fuel distribution at the end of the simulation. The black contour denotes the ignition-temperature contour defining the fire perimeter, and the red contour outlines the burned region. The white dashed line represents the containment line. A constant wind speed and direction are maintained throughout the simulation.

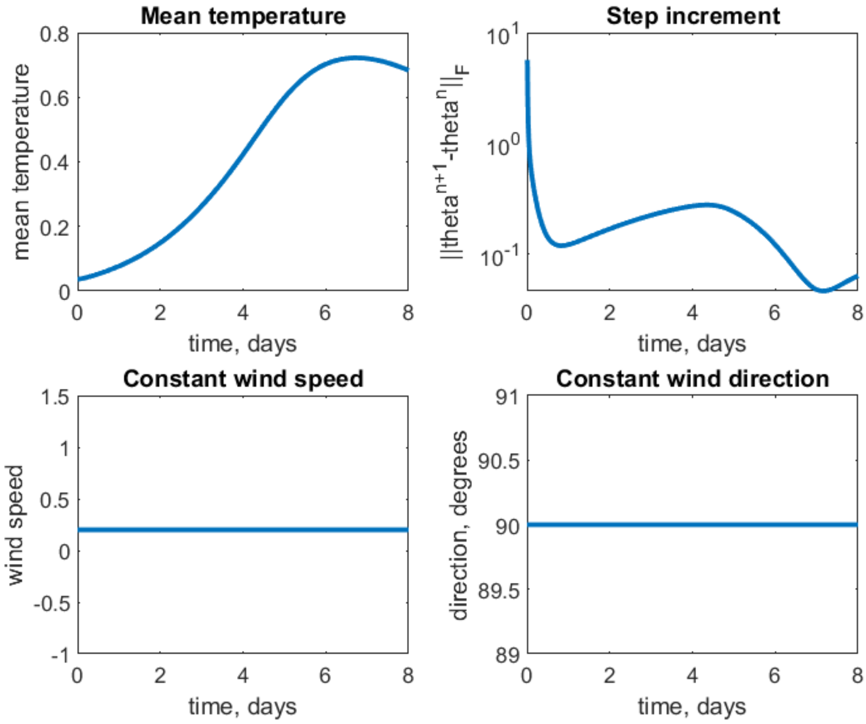

Figure 20. Diagnostic quantities for the constant-wind simulation. The panels show the temporal evolution of the spatially averaged temperature, numerical increment, wind speed, and wind direction. The mean temperature characterizes the overall thermal behavior of the wildfire, while the numerical increment provides a measure of solution convergence. Wind speed and wind direction remain constant throughout the simulation.

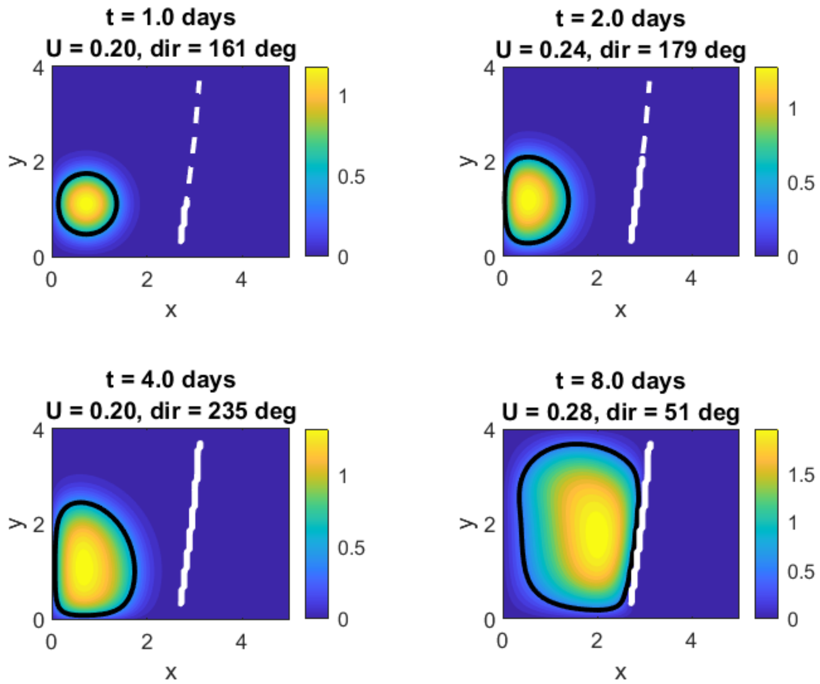

Figure 21. Temporal evolution of the wildfire temperature field subject to stochastic time-dependent wind forcing. The temperature distribution is shown at $t = 1$, 2, 4, and 8 days. Wind speed is modeled using an Ornstein–Uhlenbeck process and wind direction using a smooth correlated von Mises process. The black contour represents the ignition-temperature contour that defines the wildfire perimeter, while the white dashed line denotes the containment line. The corresponding wind speed and direction at each snapshot are indicated above the temperature fields.

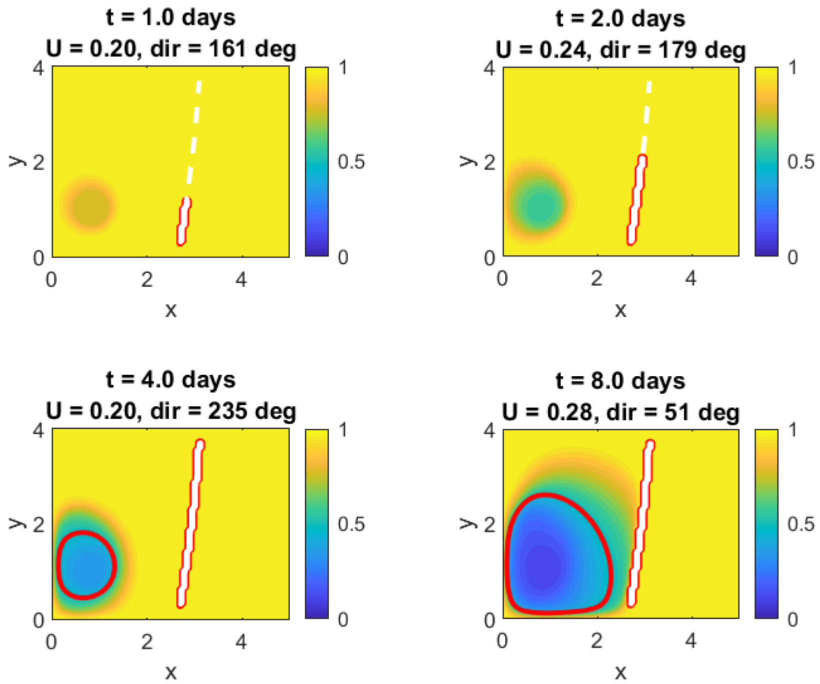

Figure 22. Temporal evolution of the wildfire fuel field subject to stochastic time-dependent wind forcing. The remaining fuel distribution is shown at $t = 1$, 2, 4, and 8 days. Wind speed is modeled using an Ornstein–Uhlenbeck process and wind direction using a smooth correlated von Mises process. The red contour outlines the burned region, while the white dashed line denotes the containment line. The corresponding wind speed and direction at each snapshot are indicated above the fuel fields.

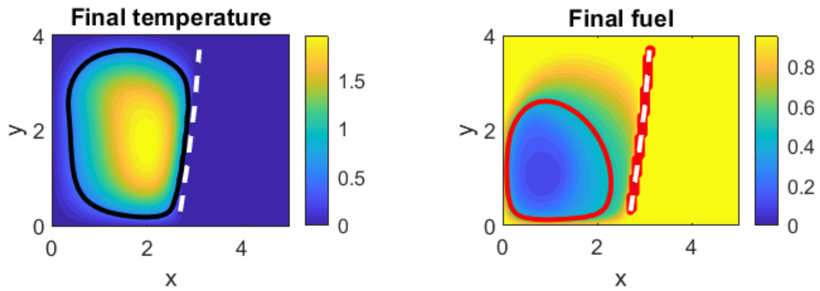

Figure 23. Final wildfire temperature and fuel distributions under stochastic time-dependent wind conditions. The left panel shows the final temperature field, while the right panel shows the remaining fuel distribution at the end of the simulation. Wind speed evolves according to an Ornstein–Uhlenbeck process and wind direction follows a smooth correlated von Mises process. The black contour in the temperature field denotes the ignition-temperature contour defining the wildfire perimeter, whereas the red contour in the fuel field outlines the burned region. The white dashed line represents the containment line. Temporal fluctuations in wind speed and direction influence the final wildfire shape and burn pattern.

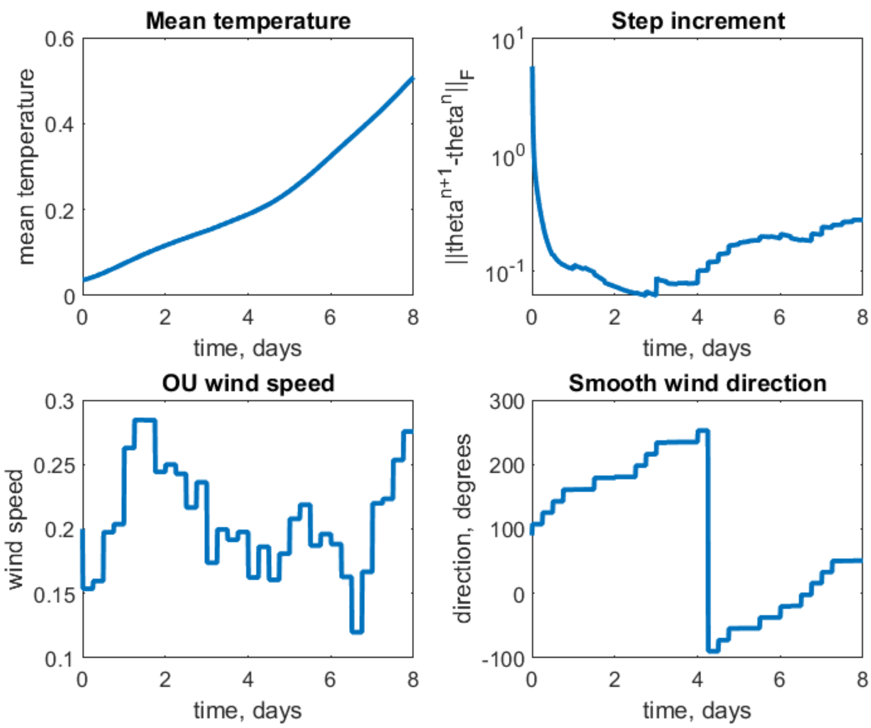

Figure 24. Diagnostic quantities for the stochastic time-dependent wind simulation. The panels show the temporal evolution of the spatially averaged temperature, numerical increment, Ornstein–Uhlenbeck (OU) wind speed, and smooth time-varying wind direction. The mean temperature characterizes the overall thermal evolution of the wildfire, while the numerical increment provides a measure of the change in the solution between successive time steps. The lower panels illustrate the stochastic fluctuations in wind speed and the corresponding variations in wind direction that influence wildfire propagation throughout the simulation.

8.3 Summary of Stochastic Wind Results

The stochastic wind simulations demonstrate that time-dependent wind forcing can substantially alter wildfire evolution compared with the constant-wind case. Under constant wind speed and direction, the fire front evolves in a smoother and more predictable manner. In contrast, the stochastic wind model introduces fluctuations in both wind speed and direction, producing a more irregular temperature field, asymmetric fuel depletion, and a different final burned region. These results demonstrate that incorporating stochastic wind forcing provides a more realistic basis for evaluating wildfire spread and the robustness of containment strategies under uncertain atmospheric conditions.

9 Conclusion

9.1 Summary of Achievements

Reflecting on the objectives outlined in Section 1, we note that we were widely successful in satisfying Objectives 1 and 2. Namely, in this report, starting with a purely diffusive toy fire-spread model, we were able to identify viable fire containment strategies/algorithms and couple these algorithms with numerical simulations, allowing for at least qualitative comparison of how well each of the strategies performed.

Having considered the simple (purely-diffusive) Fisher-KPP phenomenological model, we then progressed to a more physically motivated coupled reaction-diffusion model of fire spread, incorporating the additional physical aspects of wind and fuel depletion, thus, addressing Objective 2 outlined in Section 1. Again, the same algorithms for fire containment were considered in the context of this model.

In this report, the fire containment strategies considered utilized geometric, boundary-following and optimization-based approaches. While the models of fire spread considered were not sufficiently rich to capture the true dynamics of wildfires observed in data, the algorithms identified may easily be scaled/adapted to more realistic simulations. This suggests that Objective 4 of comparing the containment strategies identified to historical records of fire containment in practice is not out of reach.

Finally, we investigated the influence of stochastic time-dependent wind-forcing on the wildfire spread in Section 8, therefore, touching on the requirements of Objective 3. Comparisons between constant and time-dependent wind conditions demonstrated that time-dependent wind forcing can sub- substantially alter wildfire evolution, resulting in a more irregular temperature field, asymmetric fuel depletion, and a different final burned region. This suggests that incorporating randomness into our models is necessary in order to more faithfully model wildfire spread. This suggests a direction that can be taken to more realistically model wildfires, giving a stronger foundation on which we can compare the containment strategies identified.

9.2 Future Directions

Of the many directions that we could take this project, we consider three that are of particular interest.

- *Model Enhancement*: Having identified wind-randomness as having a significant impact on the spreading of the wildfire, can we incorporate other elements of uncertainty into our models? One of the most significant being fire "jump" events, where new fire nucleation sites "randomly" appear due to the carrying of embers in the wind.

- *Forecasting and Data Assimilation*: In line with Objective 3, leverage data as a means of validating simulations, and assimilate them with the theoretical models using filtering techniques to improve forecasts.
- *Robust and Probabilistic Containment*: Consider probabilistic fire containment strategies that can account for the uncertainty that is introduced when considering stochastic models of wildfire spread. Additionally, consider approaches that account for limited or uncertain knowledge of the current fire front, including utilizing a Kalman filter to combine model predictions and real-time observations.

10 Author Attributions

- MA developed and implemented the SVD algorithm in 6.5 and served as group mentor.
- NC contributed to literature review, implementing and adapting the Zambon algorithm, and writing sections 2, 6.3, A.1, 7.1
- TK contributed to Stochastic Wind Model development, implementation, numerical simulations, result analysis, figure preparation, and report writing to section 8, 2, 9 and review of the overall draft of the report.
- TN contributed to literature review, problem formulation, and analytical work, and writing sections 6.1 and 9.
- SE contributed to literature review and identified the wildfire containment model of Fried and Fried [1996] and Bressan and Wang [2010] , writing sections 4.2, 6.4, B.1, simulated and implemented Bressan & Wang "Ice cream cone" algorithm.
- TG contributed to Sections 1, 4.1, 4.3, 6.2, 6.5, 9, the Summary, and the overall structure, formatting and presentation of the report.
- VHN simulates (4.5) and implement the containment strategies, including the Fried–Fried, Zambon, geometric, and SVD-based methods and writing section 5.2.
- WT contributed to developing algorithm comparison methodology via Green's theorem and writing section 7.

11 Disclosure of AI Tool Use

- TG used AI tools to help with the generation and formatting of code, the generation of pseudocode, identify relevant literature, and debug LaTeX errors that arose with formatting this report.
- TK used AI tools to assist with code generation and debugging, report preparation. All final modeling decisions, analyses, and conclusions were reviewed and verified by the author. Any errors or omissions remain the sole responsibility of the author.
- VHN used AI as assistive tools for language editing, code generation, debugging, organization of numerical experiments, and brainstorming modeling and containment strategies. All mathematical formulation, modeling choices, numerical results, interpretation, and final content were reviewed and verified by the authors, who take full responsibility for the work.

- SE used AI in switching out codes from Matlab to Python, debugging codes and latex errors, and generation of pseudocode. Any errors and oversights are the sole responsibilities of the author.
- TN used AI tools to identify relevant literature and to locate and summarize results from identified papers, particularly when searching for a particular result that was cited in other work or had been identified by the group. All findings were manually verified to be accurate and contextually appropriate. TN also used AI to help ideate and explore possible directions for the project and the practical implications of theoretical findings, as well as explain technical concepts in conjunction with consulting other group members and traditional educational resources. Finally, TN used AI to identify and fetch real world datasets on wildfire spread. All writing, including the derivation of mathematical results in Section 6.1 was done by the authors, and where potential approaches, past results, or interpretations to include in the review arose from AI, they were evaluated and confirmed by the authors. Any errors or omissions remain the sole responsibility of the author.

References

Aishah Al Yammahi, Prashanth R. Marpu, and Taha B. M. J. Ouarda. Modeling directional distributions of wind data in the United Arab Emirates at different elevations. *Arabian Journal of Geosciences*, 14(9):774, May 2021. ISSN 1866-7511, 1866-7538. . URL https://link.springer.com/10.1007/s12517-021-06864-3.

Stefano Bianchini and Martina Zizza. Existence of spiral strategies for blocking fire spreading, August 2025. URL http://arxiv.org/abs/2508.05324. arXiv:2508.05324 [math.AP].

Alberto Bressan. Differential inclusions and the control of forest fires. *Journal of Differential Equations*, 243(2):179–207, December 2007. ISSN 0022-0396. . URL https://www.sciencedirect.com/science/article/pii/S002203960700085X.

Alberto Bressan. Dynamic Blocking Problems for a Model of Fire Propagation. In Roderick Melnik and Ilias S. Kotsireas, editors, *Advances in Applied Mathematics, Modeling, and Computational Science*, pages 11–40. Springer US, Boston, MA, 2013. ISBN 978-1-4614-5389-5. . URL https://doi.org/10.1007/978-1-4614-5389-5_2.

Alberto Bressan. A DYNAMIC PROGRAMMING APPROACH TO THE FIRE BLOCKING PROBLEM. *Communications in Optimization Theory*, 2026(23), 2026. ISSN 2051-2953. . URL https://www.scopus.com/pages/publications/105030945291.

Alberto Bressan and Maria Teresa Chiri. On the regularity of optimal dynamic blocking strategies. *Calculus of Variations and Partial Differential Equations*, 61(1):36, January 2022. ISSN 1432-0835. . URL https://doi.org/10.1007/s00526-021-02148-6.

Alberto Bressan and Tao Wang. The minimum speed for a blocking problem on the half plane. *Journal of Mathematical Analysis and Applications*, 356(1):133–144, August 2009. ISSN 0022-247X. . URL https://www.sciencedirect.com/science/article/pii/S0022247X09001802.

Alberto Bressan and Tao Wang. Equivalent formulation and numerical analysis of a fire confinement problem. *ESAIM: Control, Optimisation and Calculus of Variations*, 16 (4):974–1001, 2010. ISSN 1292-8119. URL https://eudml.org/doc/250702.

Alberto Bressan and Tao Wang. On the optimal strategy for an isotropic blocking problem. *Calculus of Variations and Partial Differential Equations*, 45(1):125–145, September 2012. ISSN 1432-0835. . URL https://doi.org/10.1007/s00526-011-0453-4.

Alberto Bressan, Maria Burago, Arthur Friend, and Jessica Jou. Blocking strategies for a fire control problem. *Analysis and Applications*, 06(03):229–246, July 2008. ISSN 0219-5305. . URL https://www.worldscientific.com/doi/abs/10.1142/S0219530508001146.

CalFire. Fire Perimeters. URL https://www.fire.ca.gov/what-we-do/fire-resource-assessment-program/fire-perimeters.

Marco Caliari and Fabio Cassini. A second order directional split exponential integrator for systems of advection–diffusion–reaction equations. *Journal of Computational Physics*, 498:112640, 2024. .

José A. Carta, Celia Bueno, and Penélope Ramírez. Statistical modelling of directional wind speeds using mixtures of von Mises distributions: Case study. *Energy Conversion and Management*, 49(5):897–907, May 2008. ISSN 01968904. . URL https://linkinghub.elsevier.com/retrieve/pii/S0196890407003767.

Yang Chen, Stijn Hantson, Niels Andela, Shane R. Coffield, Casey A. Graff, Douglas C. Morton, Lesley E. Ott, Efi Foufoula-Georgiou, Padhraic Smyth, Michael L. Goulden, and James T. Randerson. California wildfire spread derived using VIIRS satellite observations and an object-based tracking system. *Scientific Data*, 9(1):249, May 2022. ISSN 2052-4463. . URL https://www.nature.com/articles/s41597-022-01343-0.

Michael I. Davidzon. Newton's law of cooling and its interpretation. *International Journal of Heat and Mass Transfer*, 55(21-22):5397–5402, October 2012. ISSN 00179310. . URL https://linkinghub.elsevier.com/retrieve/pii/S0017931012001846.

Mark A. Finney. FARSITE: Fire Area Simulator-model development and evaluation. Technical Report RMRS-RP-4, U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station, Ft. Collins, CO, 1998. URL https://research.fs.usda.gov/treesearch/4617.

Mark A Finney. Fire growth using minimum travel time methods. *Canadian Journal of Forest Research*, 32(8):1420–1424, August 2002. ISSN 0045-5067, 1208-6037. . URL http://www.nrcresearchpress.com/doi/10.1139/x02-068.

Jeremy S. Fried and Burton D. Fried. Simulating Wildfire Containment with Realistic Tactics. *Forest Science*, 42(3):267–281, August 1996. ISSN 0015-749X. . URL https://doi.org/10.1093/forestscience/42.3.267.

Sang-Sub Kim, Rolf Klein, David Kübel, Elmar Langetepe, and Barbara Schwarzwald. Geometric firefighting in the half-plane. *Computational Geometry*, 95:101728, April 2021. ISSN 0925-7721. . URL https://www.sciencedirect.com/science/article/pii/S092577212030122X.

Rolf Klein, Christos Levcopoulos, and Andrzej Lingas. Approximation Algorithms for the Geometric Firefighter and Budget Fence Problems. *Algorithms*, 11(4):45, April 2018. ISSN 1999-4893. . URL https://www.mdpi.com/1999-4893/11/4/45.

Rolf Klein, Elmar Langetepe, Barbara Schwarzwald, Christos Levcopoulos, and Andrzej Lingas. On a Fire Fighter's Problem. *International Journal of Foundations* *of Computer Science*, 30(02):231–246, February 2019. ISSN 0129-0541. . URL https://www.worldscientific.com/doi/abs/10.1142/s0129054119500023.

Jan Mandel, Lynn S. Bennethum, Jonathan D. Beezley, Janice L. Coen, Craig C. Douglas, Minjeong Kim, and Anthony Vodacek. A wildland fire model with data assimilation. *Mathematics and Computers in Simulation*, 79(3):584–606, December 2008. ISSN 03784754. . URL https://linkinghub.elsevier.com/retrieve/pii/S0378475408001572.

Sergey Obukhov, Emad Ahmed, Denis Davydov, Talal Alharbi, Ahmed Ibrahim, and Ziad Ali. Modeling Wind Speed Based on Fractional Ornstein-Uhlenbeck Process. *Energies*, 14(17):5561, September 2021. ISSN 1996-1073. . URL https://www.mdpi.com/1996-1073/14/17/5561.

E Pastor. Mathematical models and calculation systems for the study of wildland fire behaviour. *Progress in Energy and Combustion Science*, 29(2):139–153, 2003. ISSN 03601285. . URL https://linkinghub.elsevier.com/retrieve/pii/S0360128503000170.

Mauricio J.O. Zambon, Pedro J. De Rezende, and Cid C. De Souza. Finding exact solutions for the Geometric Firefighter Problem in practice. *Computers & Operations Research*, 97:72–83, September 2018. ISSN 03050548. . URL https://linkinghub.elsevier.com/retrieve/pii/S0305054818301163.

Patrick Zippenfenig. Open-Meteo.com Weather API, December 2024. URL https://zenodo.org/doi/10.5281/zenodo.7970649.

Appendix A

A.1 Data Acquisition

All data used was taken from publicly available sources. In particular, data on historical wildfire perimeters was taken from CalFire and Chen et al. [2022], while data on wind conditions were taken from Zippenfenig [2024].

The data used primarily served to ensure our simulations had generally realistic conditions. Although we initially had ambitions to compare our wildfire spread models with historical wildfire data, we soon realized this historical data was ill-suited for this purpose; after all, most wildfires involve active human effort to contain and are therefore not a accurate representation of unchecked wildfire spread.

Appendix B

B.1 Detailed Pseudocode for FBFB Containment Algorithm

The following pseudocode gives a detailed description of the fire containment algorithm suggested by Fried and Fried [1996]. This description of the algorithm was generated with help from a Large Language Model.

Algorithm 4 Free Burning Fire Boundary Containment Algorithm (Detailed) [Fried and Fried, 1996]

Require: Initial fire size h_0, fire shape eccentricity ϵ, resource arrival time t_0, productivity rates V_i, maximum simulation time $t_{\max}$

Initialize Fire Geometry

Define the free-burning fire boundary (fbfb) as an expanding ellipse:

$$x = h(t)X(u, \epsilon), \qquad y = h(t)Y(u, \epsilon).$$

Compute head-fire spread rate

$$V_h = \frac{dh}{dt}.$$

Initialize Suppression Resources

for $i \in \{1, 2\}$ **do**

  Assign line-building rate V_i.

  Assign arrival time t_0.

  Compute efficiency ratio

$$P_i = \frac{V_i}{V_h}.$$

end for

Initialize Boundary Conditions

Set initial line-building positions u_1 and u_2 at the fire tail for direct tail attack.

$t \leftarrow t_0$

$h \leftarrow h_0$

while $t \leq t_{\max}$ **do**

  for $i \in \{1, 2\}$ **do**

Evaluate

$$\Delta_i = P_i^2(X_u^2 + Y_u^2) - (X_u Y - Y_u X)^2.$$

if $\Delta_i < 0$ **then**

return Escaped Fire

end if

Compute

$$\frac{du_i}{dh} = \frac{-(XX_u + YY_u) \pm \sqrt{\Delta_i}}{h(X_u^2 + Y_u^2)}.$$

Use the negative sign for the upper flank and the positive sign for the lower flank.

Advance u_i using a Runge–Kutta integration step.

if $\frac{du_i}{dh}$ changes sign **then**

return Escaped Fire

end if

end for

Compute containment points

$$Q_1 = (x_1, y_1), \qquad Q_2 = (x_2, y_2).$$

Compute separation distance

$$D = \|Q_1 - Q_2\|.$$

if $D = 0$ **then**

$t_c \leftarrow t$

$h_c \leftarrow h$

break

end if

Advance fire growth:

$$h \leftarrow h + \Delta h.$$

Update simulation time:

$$t \leftarrow t + \frac{\Delta h}{V_h}.$$

end while

if $t > t_{\max}$ **then**

return Escaped Fire

end if

Post-Containment Calculations

Compute containment time t_c.

Compute burned area:

$$A = \int_{h_0}^{h_c} y_1(h) \frac{dx_1}{dh}\, dh - \int_{h_0}^{h_c} y_2(h) \frac{dx_2}{dh}\, dh.$$

return $(t_c,\ A)$

B.2 Detailed Pseudocode for SVD Containment Algorithm

The following pseudocode gives a detailed description of the SVD fire containment algorithm suggested by Dr. Manuchehr Aminian. This description of the algorithm was translated from its original code using AI tools.

Algorithm 5 SVD Containment Algorithm

Require: Forecasted perimeter points $X_{\text{forecast}}(t)$
Require: Containment speed γ
Require: Number of nearest perimeter points k
Require: Desired distance from the perimeter d^*
Require: Distance-weighting falloff parameter σ
Require: Initial containment point r_0
Require: Initial direction dr
Require: Sequence of forecast times $\{t_i\}_{i=0}^{N}$

for $i = 1$ to N **do**
    $X \leftarrow X_{\text{forecast}}(t_i)$
    Select X_{near} as the k nearest points in X to r_{i-1} in Euclidean distance
    $x_{\text{bdry}} \leftarrow \text{mean}(X_{\text{near}})$
    Compute the reduced SVD of mean-subtracted neighborhood (WLOG $\sigma_1 \geq \sigma_2$)

$$X_{\text{near}} - \mathbf{1}x_{\text{bdry}}^T = USV^\top$$

    $v_\parallel \leftarrow V_1$; $v_\perp \leftarrow V_2$
    Orient the parallel vector consistent with previous motion:

$$p \leftarrow \text{sign}(v_\parallel \cdot dr)\, v_\parallel$$

    Orient the perpendicular direction relative to the nearby boundary:

$$q \leftarrow -\,\text{sign}\left(v_\perp \cdot (r_{i-1} - x_{\text{bdry}})\right) v_\perp$$

    $d \leftarrow \|r_{i-1} - x_{\text{bdry}}\|$
    Compute the distance-based weighting for $v_\parallel$ and $v_\perp$:

$$\alpha \leftarrow \exp\left(\frac{-(d-d^*)^2}{\sigma^2}\right)$$

    Compute the containment direction:

$$w \leftarrow \alpha p + \sqrt{1-\alpha^2}\, q$$

    Advance the containment distance γ:

$$r_i \leftarrow r_{i-1} + \gamma w$$

    $dr \leftarrow w$
end for
return containment path $\{r_i\}_{i=0}^{N}$