Differentiable Tactical Optimization of Aerial Wildfire Suppression

Ion Matei
Fujitsu Research of America
imatei@fujitsu.com

Maksym Zhenirovskyy
Fujitsu Research of America
mzhenirovskyy@fujitsu.com

Takuya Kurihana
Fujitsu Research of America
tkurihana@fujitsu.com

Rohit Vuppala
Fujitsu Research of America
rvuppala@fujitsu.com

Anthony Wong
Fujitsu Research of America
awong@fujitsu.com

Abstract

We present a differentiable framework for optimizing aerial wildfire suppression and evaluating its robustness under uncertainty. The method couples a pretrained hybrid convolutional neural network–cellular automaton (CNN–CA) wildfire simulator with a gradient-based optimizer over binary aerial drops with continuous timing, location, and orientation. The intervention model distinguishes between water, which suppresses currently burning cells, and retardant, which persistently reduces future spread through a fuel-suppression mechanism. We quantify both aleatoric uncertainty, via Monte Carlo sampling of daily fire-state realizations, and epistemic uncertainty, via spatially correlated prediction-error perturbations in compositional state space. Experiments on the 2020 Bear Fire show that the framework can generate effective intervention schedules for total fire-area minimization while providing uncertainty-aware estimates of suppression benefit. Across the aleatoric, epistemic, and combined uncertainty evaluations, the optimized interventions reduce the final fire-affected area by an average of 92.8%.

1 Introduction

Wildfire suppression is a difficult sequential decision problem. Fire spread depends on terrain, fuels, weather, and the current fire state, while suppression actions must be selected quickly and under uncertainty. The challenge is not only to predict fire evolution, but also to decide when, where, and how to intervene. Wildfire modeling has progressed from semi-empirical and rule-based formulations [15], to cellular-automata models with explicit local spread rules [1, 7], to machine-learning and hybrid approaches that adapt spread behavior to data [10, 19]. In parallel, wildfire suppression planning has been studied through operations-research models for resource allocation [2, 3, 8], strategic simulation frameworks [9, 14], and predictive tools for identifying control opportunities or attack success [6, 13]. However, these lines of work are usually separated. Existing methods rarely combine a learned spatial fire model, direct optimization of intervention geometry, and explicit uncertainty quantification in one framework. We address this gap with a unified framework for tactical aerial wildfire intervention design. The method uses a pretrained hybrid CNN–CA simulator that predicts spatially varying fire dynamics from terrain, fuel, and weather inputs, and then optimizes intervention schedules by differentiating through the simulator. The optimizer acts on binary aerial drops with continuous spatial geometry, while aircraft availability and grounded-day constraints are enforced during rollout. The method also distinguishes between water and retardant: water acts immediately on active burning, whereas retardant persistently suppresses future spread. After optimization, the same framework evaluates the resulting intervention under both stochastic

2026 Conference on Physics and AI (PAI 2026).

spread variability and learned model discrepancy. Our main contributions are: (i) a differentiable formulation of aerial wildfire intervention design based on a hybrid CNN–CA model, with binary drop decisions, continuous geometry, and distinct physical semantics for water and retardant; and (ii) an uncertainty-aware evaluation framework that quantifies both aleatoric and epistemic uncertainty in intervention outcomes. Our approach belongs to the emerging class of differentiable and hybrid wildfire simulators [4, 16, 18, 19], but extends this line of work from forward prediction to tactical intervention design under uncertainty.

2 Method

At time t and raster cell $x \in \Omega$, the wildfire state is represented by a three-state probability vector $\mathbf{p}_t(x) = \big(p_t^U(x), p_t^B(x), p_t^R(x)\big)$, where U, B, and R denote unburned, burning, and burned. The fire-occupancy probability is $P_t^{\text{fire}}(x) = p_t^B(x) + p_t^R(x) = 1 - p_t^U(x)$. The wildfire dynamics are generated by a frozen hybrid CNN–CA model. The proposed optimization framework is not tied to a particular training regime for the wildfire simulator. A simulator trained across multiple fire events could potentially improve predictive accuracy and generalization, but it would not change the structure of the proposed intervention-optimization methodology. The CNN predicts spatially varying propagation parameters from terrain, fuel, and wind inputs, and the CA propagates the probabilistic fire state through local wind-, slope-, and fuel-modulated interactions. On top of this simulator, we optimize aerial intervention schedules over aircraft-time binary drop variables. Other optimization variables include the continuous pose parameters that determine the location and orientation of each drop footprint. Each intervention is modeled as a wind-corrected anisotropic footprint, centered at the predicted landing point. The footprint is represented by an elliptical Gaussian deposit whose center and orientation are determined by the optimized pose, and whose principal axes depend on aircraft speed, release duration, drop height, and wind-induced drift. The deposited suppression field depends on aircraft payload, footprint geometry, and material type. Water acts immediately by reducing the current burning-state probability, whereas retardant induces a persistent multiplicative reduction of the effective fuel term used by the CA at future steps. Binary drop execution is handled with a straight-through estimator so that the forward simulation uses discrete actions while gradients remain available during backpropagation. Aircraft turnaround and grounded-day constraints are enforced during rollout, so the optimizer searches only over feasible intervention schedules. The optimization objective combines fire exposure, terminal fire extent, intervention cost, and optional protected-region penalties: $\mathcal{L} = \lambda_{\text{burn}}\mathcal{L}_{\text{burn}} + \lambda_{\text{final}}\mathcal{L}_{\text{final}} + \lambda_{\text{budget}}\mathcal{L}_{\text{budget}}$, with $\mathcal{L}_{\text{burn}} = \frac{1}{T|\Omega|}\sum_{t=1}^{T}\sum_{x\in\Omega} P_t^{\text{fire}}(x)$, $\mathcal{L}_{\text{final}} = \frac{1}{|\Omega|}\sum_{x\in\Omega} P_T^{\text{fire}}(x)$, $\mathcal{L}_{\text{budget}} = \sum_{t,a} d_{t,a}$. Here $d_{t,a} \in \{0,1\}$ is the executed drop decision for aircraft a at time t. We optimize the intervention variables with Adam and then apply a second-stage budget-refinement pass that removes unnecessary sorties while penalizing deviation from the fire-performance level achieved by the reference solution.

3 Uncertainty-Aware Evaluation

The deterministic rollout yields a single wildfire trajectory, but operational use requires uncertainty-aware performance estimates. We therefore consider two complementary uncertainty sources. **Aleatoric uncertainty:** We sample daily binary fire-state realizations from the model-implied end-of-day categorical distribution while keeping the within-day CA evolution deterministic. This preserves the learned within-day accumulation behavior of the simulator while producing stochastic day-to-day wildfire trajectories. **Epistemic uncertainty:** We model prediction error as a spatially correlated random field in isometric log-ratio (ILR) coordinates. The ILR transform maps each 3-state probability vector from the simplex to $\mathbb{R}^2$, allowing additive Gaussian perturbations that preserve validity after inverse mapping. The covariance and spatial correlation length of the perturbation are learned from held-out residuals between deterministic predictions and observed fire states. These uncertainty models can be applied to both the baseline and the intervention rollout. In the intervention case, the optimized schedule is held fixed, so uncertainty is quantified *conditional on the chosen policy*. This makes it possible to evaluate not only nominal effectiveness, but also its variability under stochastic spread and model discrepancy.

Table 1: Uncertainty-aware intervention performance on the 2020 Bear Fire. Total burned area means (in *hectares*) are reported at the final horizon. Reduction is relative to the corresponding no-intervention baseline.

UQ type	Baseline	Optimized	Reduction
Aleatoric	9093.6	330.1	96.4%
Epistemic	12376.6	1882.5	84.8%
Combined	6950.1	204.9	97.1%

4 Experiments

We evaluate the framework on a suppression scenario based on the 2020 Bear Fire event. This experimental setting should be interpreted as a proof-of-concept evaluation of differentiable tactical intervention optimization, rather than as a claim about cross-event wildfire-model generalization. The hybrid CNN–CA wildfire model was pretrained on the same 2020 Bear Fire dataset used for the present study. The wildfire simulator uses static terrain and fuel information from LANDFIRE, daily meteorological forcing from ECMWF ERA5, and daily observed fire perimeters from the dataset of [17]. We optimize for total-fire-area minimization using 21 fixed-wing firefighting aircraft, including S-2T and AT-802F tactical platforms, BAe-146/RJ85, MD-87, C-130 MAFFS, DC-10, one Boeing 747 Supertanker, and CL-415 Super Scoopers. The interventions are optimized over a 15-day event horizon, with grounded-day constraints reflecting the temporary suspension of fixed-wing operations on September 9–10, 2020. The fleet and grounded-day constraints are chosen to reflect the tiered aerial response and operational disruptions observed during the Bear Fire [5, 11, 12]. The optimizer first searches for a strong suppression policy and then performs a budget-refinement step to reduce resource usage while preserving fire-performance targets. Figure 1 compares the final fire-occupancy maps for the baseline and optimized cases, Figure 2 shows the resulting spatial distribution of optimized aerial drops, and Figure 3 reports the uncertainty-aware baseline and intervention trajectories under aleatoric and epistemic uncertainty. Table 1 summarizes the corresponding final-horizon performance under aleatoric, epistemic, and combined uncertainty. Figure 1 shows that the optimized schedule substantially reduces the final fire footprint relative to the no-intervention baseline. The corresponding drop pattern in Figure 2 indicates that the optimizer concentrates resources near locations that most influence subsequent spread, rather than distributing drops uniformly over the active fire. The uncertainty results in Figure 3 and Table 1 show that the intervention remains beneficial under aleatoric, epistemic, and combined uncertainty. The epistemic trajectories are broader and less optimistic because they reflect structured, spatially correlated model error. The aleatoric trajectories can appear more optimistic than the deterministic rollout because daily binary sampling removes weak residual burning mass that would otherwise continue to influence spread in the mean-field simulator. Relative to prior wildfire-suppression methods, our approach occupies a different point in the design space. Operations-research formulations typically focus on allocation, routing, or containment decisions under prescribed fire dynamics [2, 3, 8]. Strategic planning frameworks emphasize long-horizon scenario comparison [9, 14]. Predictive support tools estimate likely control opportunities or attack success [6, 13]. In contrast, our method directly differentiates through a learned wildfire simulator to produce a time-indexed aerial intervention policy together with an uncertainty-aware evaluation of its outcomes.

5 Conclusion

We presented a differentiable framework for aerial wildfire suppression optimization under uncertainty. The method combines a pretrained hybrid CNN–CA wildfire model, gradient-based optimization of binary drops with continuous geometry, physically distinct water and retardant effects, and post hoc uncertainty quantification through aleatoric and epistemic models. Results on the 2020 Bear Fire show that the approach can generate effective intervention schedules for total fire-area minimization while also providing uncertainty-aware estimates of suppression benefit. Several extensions are left for future work. First, future studies will test the method under simulators trained across multiple events and regions. Second, we will perform ablation studies against non-differentiable and rule-based intervention strategies, such as perimeter buffering, uniform resource allocation, or greedy drop placement.

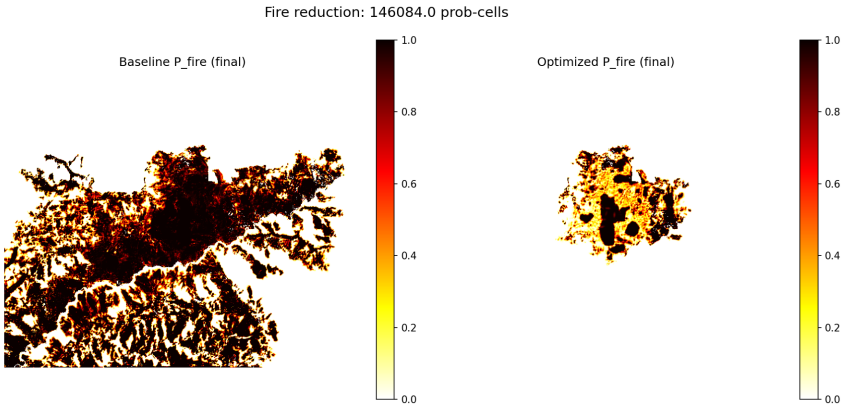

Figure 1: Total fire-area minimization: final fire-occupancy maps for the baseline and optimized intervention cases.

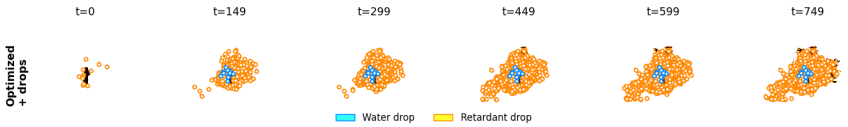

Figure 2: Total fire-area minimization: spatial distribution of the optimized aerial drops over time.

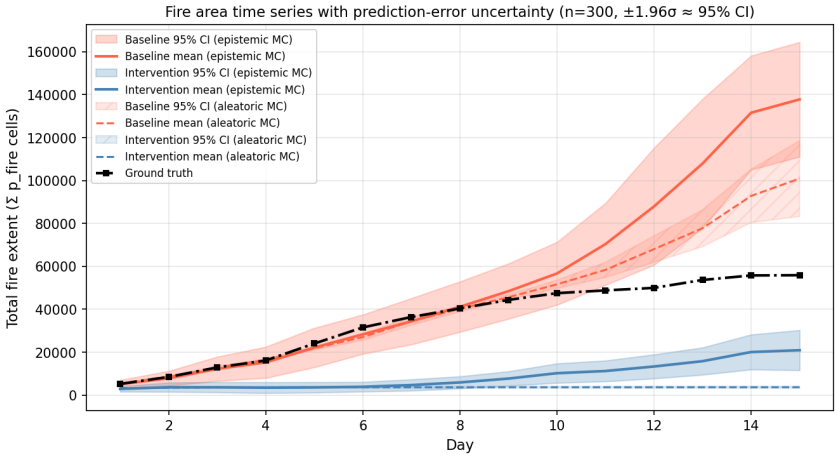

Figure 3: Total fire-area minimization: baseline and intervention trajectories for total fire-affected area under aleatoric and epistemic uncertainty quantification.

References

[1] A. Alexandridis, D. Vakalis, C.I. Siettos, and G.V. Bafas. A cellular automata model for forest fire spread prediction: The case of the wildfire that swept through Spetses island in 1990. *Applied Mathematics and Computation*, 204(1):191–201, 2008.

[2] Mualla Gonca Avci, Mustafa Avci, Maria Battarra, and Güneş Erdoğan. The wildfire suppression problem with multiple types of resources. *European Journal of Operational Research*, 316:488–502, 2024.

[3] Erin J. Belval, Yu Wei, and Michael Bevers. A mixed integer program to model spatial wildfire behavior and suppression placement decisions. *Canadian Journal of Forest Research*, 45:384–393, 2015.

[4] Ufuk Çakır, Victor-Alexandru Darvariu, Bruno Lacerda, and Nick Hawes. JaxWildfire: A GPU-accelerated wildfire simulator for reinforcement learning. *arXiv preprint arXiv:2512.06102*, 2025.

[5] California Department of Forestry and Fire Protection (CAL FIRE). 2020 Fire Siege Report. Technical report, State of California, 2021. Detailed analysis of North Complex and Bear Fire operations.

[6] Adrián Cardil, Adrián Jiménez-Ruano, Santiago Monedero, Phillip SeLegue, Macarena Ortega, Raúl Quilez, Jeff Fuentes, Geoff Marshall, Robert Clark, Tim Chavez, Sergio de Miguel, Tomás Quiñones, Fernando Bezares, Scott Purdy, Kristen Allison, Tim Metzger, Chris Waters, Midhun Mohan, and Joaquín Ramirez. Assessing the suppression difficulty of wildland fires for initial attack response. *International Journal of Wildland Fire*, 34:WF24160, 2025.

[7] José G. Freire and Carlos C. DaCamara. Using cellular automata to simulate wildfire propagation and to assist in fire management. *Natural Hazards and Earth System Sciences*, 19(1):169–179, 2019.

[8] Bibiana Granda, Javier León, Begoña Vitoriano, and John Hearne. Decision support models and methodologies for fire suppression. *Fire*, 6(2), 2023.

[9] Rachel M. Houtman, Claire A. Montgomery, Aaron R. Gagnon, David E. Calkin, Thomas G. Dietterich, Sean McGregor, and Mark Crowley. Allowing a wildfire to burn: Estimating the effect on future fire suppression costs. *International Journal of Wildland Fire*, 22:871–882, 2013.

[10] Piyush Jain, Sean C.P. Coogan, Sriram G. Subramanian, Mark Crowley, Steve Taylor, and Mike D. Flannigan. A review of machine learning applications in wildfire science and management. *Environmental Reviews*, 28(4):478–505, 2020.

[11] National Interagency Coordination Center (NICC). 2020 Wildland Fire Summary and Statistics Annual Report. Technical report, National Interagency Fire Center, Boise, ID, 2021.

[12] National Wildfire Coordinating Group (NWCG). Incident Status Summary (ICS-209): North Complex CA-PNF-001308. Technical report, September 2020. Official daily operational logs for the Bear Fire period.

[13] Christopher D. O'Connor, David E. Calkin, and Matthew P. Thompson. An empirical machine learning method for predicting potential fire control locations for pre-fire planning and operational fire management. *International Journal of Wildland Fire*, 26:587–597, 2017.

[14] Karin L. Riley, Matthew P. Thompson, Joe H. Scott, and Julie W. Gilbertson-Day. A model-based framework to evaluate alternative wildfire suppression strategies. *Resources*, 7(1), 2018.

[15] Richard C Rothermel. A mathematical model for predicting fire spread in wildland fuels. Research Paper INT-115, USDA Forest Service, Intermountain Forest and Range Experiment Station, 1972.

[16] Chaopeng Shen, Alison P. Appling, Pierre Gentine, Toshiyuki Bandai, Hoshin Gupta, Alexandre Tartakovsky, et al. Differentiable modelling to unify machine learning and physical models for geosciences. *Nature Reviews Earth & Environment*, 4(8):552–567, 2023.

[17] Zeyu Xia and Sibo Cheng. Data for: "PyTorchFire: A GPU-accelerated wildfire simulator with differentiable cellular automata". Mendeley Data, V1, 2024. Dataset.

[18] Zeyu Xia and Sibo Cheng. PyTorchFire: A GPU-accelerated wildfire simulator with differentiable cellular automata. *Environmental Modelling & Software*, 188:106401, 2025.

[19] Zhong Zheng, Wei Huang, Songnian Li, and Yongnian Zeng. Forest fire spread simulating model using cellular automaton with extreme learning machine. *Ecological Modelling*, 348:33–43, 2017.