Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# Wildfire prediction for resource prioritisation

**Joe McNorton**

joe.mcnorton@ecmwf.int

European Centre for Medium-Range Weather Forecasts

**Francesca Di Giuseppe**

European Centre for Medium-Range Weather Forecasts

**Article**

**Keywords:**

**Posted Date:** November 23rd, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-11619907/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** No competing interests reported.

# Wildfire prediction for resource prioritisation

Joe McNorton[1]* and Francesca Di Giuseppe[1]

[1]European Centre for Medium-Range Weather Forecasts, Reading, UK RG2 9AX

*Corresponding author e-mail: joe.mcnorton@ecmwf.int

**Abstract**

Wildfire management agencies must allocate limited resources, particularly when concurrent fires compete for shared assets. Forecasts such as the Probability of Fire (PoF) predict where fires are likely to occur, but not their potential size. Here, we combine fire probabilities with historical and forecast burned area to prioritise locations where extensive burning is most likely. Across the Mediterranean, Canada and Australia during 2020–2025, we optimise this combination by evaluating 197 candidate scores based on their capture of total burned area and extreme fires (with BA ≥10 $km^2$ per cell-day), assessing decision value using a capacity-constrained cost–loss framework. When attention was limited to 0.1% of each domain, incorporating burned area information captured 1.89, 1.23, and 2.77 times as many extreme events as PoF alone in the Mediterranean, Canada, and Australia, respectively, while capturing 8%, 14%, 40% more total burned area. The optimal combined score varied by region, with historical burned area adding value in the Mediterranean, where fires recur in similar landscapes, but reducing performance in Canada. Under cost and loss assumptions, the optimal score resulted in the highest or near-highest net economic benefit in all domains and remained viable at higher action costs than PoF alone. Combining fire likelihood with potential extent therefore improves prioritisation under limited capacity, providing a framework adaptable to regional costs, losses and consequence metrics and offers a practical basis for forecast-informed fire management.

## Introduction

Concurrent wildfires can place competing demands on finite fire-management resources. The 2019–2020 Australian fire season, successive severe seasons in Canada during 2023–2025, and the 2025 and 2026 European fire seasons illustrate how widespread fire activity can strain emergency response capacity, disrupt communities, and generate substantial economic losses[1–4]. When fires develop simultaneously across neighbouring regions, demand for shared aircraft, crews, equipment, and situational-awareness capacity can exceed availability[5,6]. Effective preparedness therefore depends on anticipating both where fires are likely to occur and where they could burn extensively.

Forecasts of fire occurrence help identify locations that warrant increased attention. Forecast systems such as the Probability of Fire (PoF) use weather, vegetation, and

human influences to produce skillful daily predictions across large spatial domains[7,8]. However, PoF models predict fire occurrence without directly quantifying the potential extent of burning. Prioritisation based on occurrence alone may therefore favour a location likely to experience a small fire over one where fire is less likely but could, conditional on occurrence, burn extensively.

Burned area provides complementary information on fire magnitude. Extensive burning can increase suppression demands, emissions, smoke exposure, and ecological and economic impacts, although these consequences also depend on where and how fires burn[9–12]. Predictions of potential burned area may therefore help distinguish between locations with similar fire probabilities. However, a large predicted burned area conditional on fire occurrence does not necessarily imply that fire is likely. Its value for immediate prioritisation consequently depends on how it is combined with fire-occurrence probability.

Historical observations and daily predictions capture different aspects of potential fire extent. A gridded burned area climatology characterises the typical magnitude of past fires at each location, whereas daily conditional burned area forecasts reflect contemporaneous environmental conditions. Their relative value may depend on the recurrence of burning in similar landscapes and on the extent to which daily conditions determine fire size. There is therefore no single obvious way to combine these sources. A priority score must determine their relative weighting, whether burned area information is applied across the full domain or only where PoF is high, and how the resulting ranking performs at the action capacity available to decision makers[13].

Here, we evaluate whether and how burned area information improves fire management prioritisation across the Mediterranean, Canada, and Australia during 2020–2025. We compare a PoF only score with 197 candidate scores that incorporate fixed conditional burned area climatology, daily conditional burned area forecasts (FBA), or both. Each score ranks grid cells within a fixed daily action capacity, representing potential targets for enhanced monitoring, preparedness, resource pre-positioning, or situational awareness. We evaluate policies by the percentage of total observed burned area captured, the percentage of extreme events captured, where an extreme event is a gridcell day with ≥10 $km^2$ burned, and the percentage of all events captured.

We examine whether the selected score for a given capacity differs among regions, and whether improved prioritisation can provide greater potential decision value under assumed action costs and avoided losses. The cost–loss analysis is intended as a scenario-based assessment rather than an estimate of realised operational costs or economic damages[14]. Burned area is used as a measurable proxy for fire magnitude and potential loss, but the framework can accommodate alternative targets, including fire intensity, rate of spread, smoke emissions, exposure, suppression difficulty, and direct economic damages.

## Results

We evaluated whether combining daily PoF predictions with conditional burned area information improved the prioritisation of limited fire-management attention across the Mediterranean, Canada, and Australia during 2020–2025. Combining these sources requires decisions about the relative weighting of fire likelihood and potential fire extent, and whether burned area information should influence all cells or only those with high PoF.

At a daily action capacity of 0.1% of the complete domain grid, we ranked 197 candidate priority scores independently in each region by the equally weighted mean of total observed burned area capture and extreme-event capture. A full list of candidate scores is provided in the Supplementary Information. Extreme events were defined as gridcell days with at least 10 $km^2$ burned.

The selected score differed among regions (Table 1). In the Mediterranean, the highest combined capture was obtained using a PoF-gated climatology–prediction formulation (10.1%). This score combined PoF with fixed conditional burned area climatology and applied a bounded FBA adjustment only within the highest 1% of grid cells ranked by PoF. Within this gate, $\lambda = 0.8$ controlled the magnitude of the adjustment, allowing the FBA-based multiplier to range from 0.2 to 1.8 according to normalised FBA. Full formulation details are provided in the Methods. At the primary capacity, it captured 14.1% of observed burned area and 6.1% of extreme events. In Canada and Australia, the selected score was the direct product of PoF and FBA (PoF × FBA). It captured 9.8% of total observed burned area and 14.2% of extreme events in Canada, giving a combined capture of 12.0%. In Australia, it captured 3.8% of both total observed burned area and extreme events, giving a combined capture of 3.8%.

These differences indicate that the value of conditional burned area information was region dependent. In Canada and Australia, FBA performed best when multiplied directly by PoF. In the Mediterranean, performance improved when the broader PoF × climatology ranking was retained and FBA was used only to refine priorities among locations with high fire occurrence probability. We then compared the selected regional scores with PoF only, PoF × burned area climatology, and random selection across a range of daily action capacities. This comparison tested whether incorporating burned area information enabled the same limited capacity to capture more total burned area and extreme events, even if it captured fewer burned gridcell days overall. Finally, we assessed whether this shift towards larger events translated into greater decision value using a cost–loss analysis.

*Table 1. Best-performing candidate from each priority-score family in Australia, Canada, and the Mediterranean at a daily action capacity of 0.1% of the complete study-domain grid. Candidate families were evaluated independently within each region during 2020–2025 and ranked by the equally weighted mean of total observed burned-area capture and extreme event capture. Extreme events were defined as cell-days with observed burned area ≥10 $km^2$. The highest combined capture within each region is shown in bold.*

| Score Family | Configuration | Total Burned Area Captured (%) | Extreme Event Captured (%) | Combined Capture (%) |
|---|---|---|---|---|
| **Australia** | | | | |

| | | | | |
|---|---|---|---|---|
| **Multiplicative** | **PoF × FBA** | **3.825** | **3.821** | **3.823** |
| Bounded residual | λ = 1.0 | 3.149 | 3.482 | 3.316 |
| PoF-gated FBA adjustment | 2% gate, λ = 1.0 | 3.166 | 3.42 | 3.293 |
| Two-stage screening | PoF 2%, λ = 1.0 | 2.401 | 3.256 | 2.828 |
| Percentile-rank blend | 80%/15%/5% | 2.889 | 3.669 | 3.279 |
| Geometric blend | $PoF^{0.75} \times clim^{1.0} \times FBA^{0.1}$ | 2.919 | 1.747 | 2.333 |
| PoF with weighted FBA | PoF 1% - 50% FBA | 2.927 | 3.429 | 3.178 |
| **Canada** | | | | |
| **Multiplicative** | **PoF × FBA** | **9.834** | **14.160** | **11.997** |
| Bounded residual | λ = 0.9 | 3.11 | 2.805 | 2.957 |
| PoF-gated FBA adjustment | 2% gate, λ = 1.0 | 3.25 | 2.704 | 2.977 |
| Two-stage screening | PoF 0.2%, λ = 0.4 | 6.305 | 6.489 | 6.397 |
| Percentile-rank blend | 80%/15%/5% | 3.755 | 2.737 | 3.246 |
| Geometric blend | $PoF^{1.0} \times clim^{0.75} \times FBA^{0.1}$ | 7.422 | 8.584 | 8.003 |
| PoF with weighted FBA | 50% FBA | 8.173 | 10.206 | 9.190 |
| **Mediterranean** | | | | |
| Multiplicative | PoF × FBA | 13.301 | 3.964 | 8.632 |
| Bounded residual | λ = 0.6 | 13.750 | 6.126 | 9.938 |
| **PoF-gated FBA adjustment** | 1% gate, λ = 0.8 | **14.103** | **6.126** | **10.114** |
| Two-stage screening | PoF 0.2%, λ = 0.7 | 12.810 | 3.964 | 8.387 |
| Percentile-rank blend | 70%/20%/10% | 10.847 | 8.829 | 9.838 |
| Geometric blend | $PoF^{0.75} \times clim^{1.0} \times FBA^{0.05}$ | 13.546 | 3.784 | 8.665 |
| PoF with weighted FBA | PoF 1% - 40% FBA | 10.247 | 7.387 | 8.817 |

## Mediterranean Prioritisation Performance

At the primary daily action capacity of 0.1% of the complete Mediterranean grid, corresponding to 162 selected grid cells per day (~10,000 $km^2$), PoF only prioritisation captured 13.0% of total observed burned area, 3.2% of extreme events, and 12.8% of all events during 2020–2025 (Table 2). Random selection represents a system with no skill.

The selected Mediterranean formulation captured 14.1% of total observed burned area and 6.1% of extreme events. Compared with PoF only prioritisation, this represented relative improvements of 8.0% for total burned area capture and 88.9% for extreme event capture. Although, the selected score captured 12.2% of burn events, compared with 12.8% for PoF only prioritisation, a relative reduction of 5.1%.

The burned area climatology accounted for a considerable fraction of the improvement relative to PoF alone. At 0.1% capacity, PoF × fixed climatology captured 13.6% of total observed burned area, accounting for around half of the improvement observed in the selected score. The contribution from the climatology was even larger when considering the selection of extreme events, although selected score still further improved on merely using the climatology alone.

The advantage of the climatology- and FBA-informed scores over PoF only prioritisation persisted across larger action capacities (Figure 1a–c). At a daily capacity of 5% of the complete grid, PoF only prioritisation captured 57.5% of total observed burned area, 40.0% of extreme events, and 60.8% of all burn events. The selected score captured 64.2%, 52.8%, and 63.0%, respectively, corresponding to relative improvements of 11.7%, 32.0%, and 3.5%. At this capacity, the selected score and PoF × climatology showed near-identical performance, indicating that the principal benefit at larger

capacities arose from incorporating the fixed climatological signal. This is consistent with the PoF-gated design, in which the FBA adjustment is applied only within the highest 1% of daily PoF values. As the optimal score was selected at the 0.1% capacity, its performance should be interpreted principally in this low-capacity setting; alternative FBA gates or score configurations could potentially be optimised for higher action capacities. Notably, the selected score captured more burn events than PoF at the 5% capacity, despite PoF only being designed to predict fire occurrence. This suggests that, at broader action capacities, the climatological burned-area component also improves coverage of locations with recurrent observed burning.

The spatial components of the multiple scores are illustrated for 17 September 2024 (Figure 1d–g), when substantial observed burned area clusters occurred in north-western Iberia and eastern Europe (Figure 1d). The PoF prediction identified several areas with elevated likelihood of fire occurrence (Figure 1e), while FBA provided additional information on the predicted magnitude of burning in the event of a fire (Figure 1f). By combining these signals with the fixed burned area climatology, the selected score ranking concentrated the highest-priority cells in areas associated with the largest observed fire clusters (Figure 1g).

*Table 2. Mediterranean score performance at the daily action capacity of 0.1% of the complete domain during 2020–2025, corresponding to a mean of 162 selected grid cells per day. Values show the total percentage of total observed burned area, extreme events, and number of events captured by each prioritisation score. Random selection denotes the expected capture under an unskilled score.*

| Priority Score (0.1% capacity) | Total Burned Area Capture (%) | Extreme Event Capture (%) | Number of Events Captured (%) |
|---|---|---|---|
| PoF only | 13.05 | 3.24 | 12.82 |
| PoF × climatology | 13.57 | 5.77 | 11.59 |
| PoF-gated FBA adjustment | 14.10 | 6.13 | 12.17 |
| Random Selection | 0.1 | 0.1 | 0.1 |

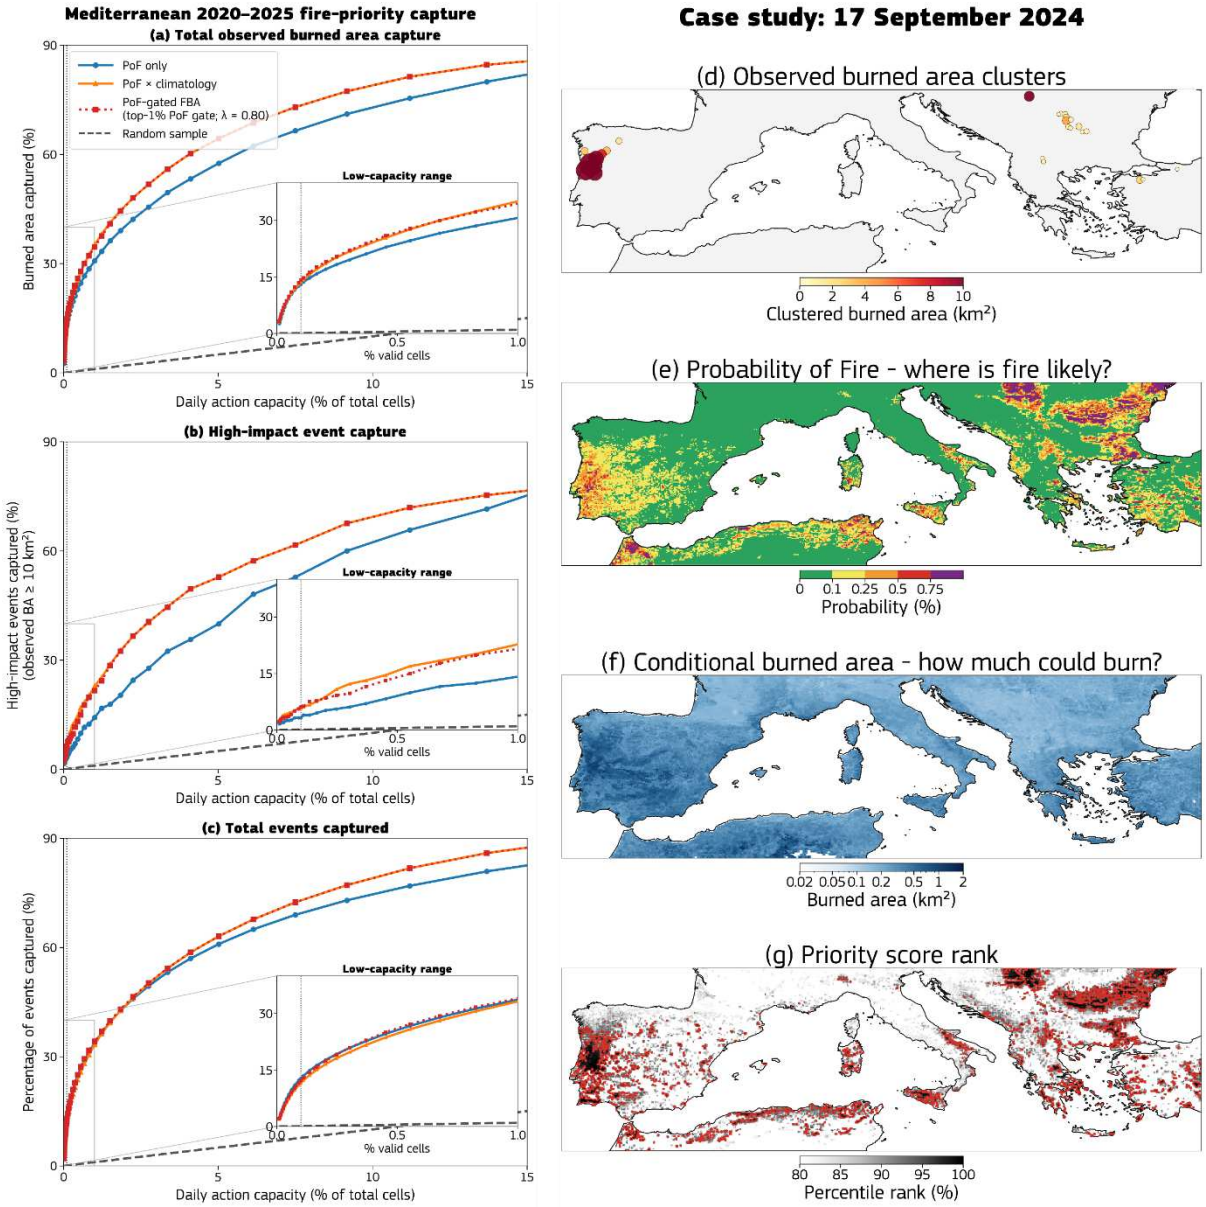



*Figure 1.* ***Mediterranean fire priority evaluation and illustrative case study.*** *(a–c) 2020–2025 capture curves for (a) total observed burned area, (b) extreme events (≥10km²/cell/day) and (c) number of gridcells with fire. Curves show the percentage captured as daily action capacity increases. Blue indicates Probability of Fire (PoF) only prioritisation; orange indicates PoF × burned area climatology from 2003-2019; red indicates the selected score, which applies a bounded conditional burned area prediction (FBA) adjustment within the daily top 1% of PoF-ranked cells (λ = 0.80); and dashed grey indicates random selection. Insets show the low-capacity range of up to 1%. (d–g) Example daily outputs for 17 September 2024 for (d) clustered observed burned area, (e) PoF predictions, (f) FBA, and (g) PoF-gated FBA adjustment percentile-ranked priority score. In panel (g), red outlines identify contiguous regions whose scores are ≥95th percentile as an illustrative example of priority regions. The example date was selected based on the number of events, total burned area and number of extreme events.*

## Australia Prioritisation Performance

For the Australian domain, an action capacity of 0.1% corresponds to 291 grid cells per day. PoF only prioritisation captured 2.7% of total observed burned area, 1.4% of extreme events, and 3.7% of all burn events during 2020–2025 (Table 3), considerably lower than in the Mediterranean. For total burned area and burn events, this difference largely reflects the much greater volume of burning relative to the selection budget. A perfect prediction could have captured at most 30.6% of events and 77.9% of total burned area,

compared with 82.2% and 97.0% in the Mediterranean. In contrast, a perfect prediction could have captured 96.8% of Australian extreme events, indicating that the low extreme-event capture reflects the limited ability of occurrence probability alone to identify where large fires occur.

The selected Australian formulation, PoF × FBA, captured 3.8% of total observed burned area and 3.8% of extreme events. Compared with PoF only prioritisation, this represented relative improvements of 39.8% for total burned area capture and 177.3% for extreme event capture, nearly tripling the number of extreme events captured. The selected score captured 3.2% of burn events, compared with 3.7% for PoF only prioritisation, a relative reduction of 13.5%. This trade-off is consistent with the candidate selection objective.

In contrast to the Mediterranean, the daily FBA provided most of the improvement over PoF alone. The selected score does not include the fixed climatological term, and PoF × fixed climatology captured 3.0% of total observed burned area and 3.0% of extreme events at 0.1% capacity. The climatology therefore recovered only around 21% of the burned area improvement and 68% of the extreme event improvement achieved by the multiplicative score. PoF × fixed climatology also captured the fewest burn events of the three scores (2.1%). These results indicate that in Australia the daily FBA carries information on burning magnitude beyond the recurrent spatial pattern represented by the climatology.

The advantage of the selected multiplicative score over PoF only prioritisation persisted across larger action capacities (Figure 2a–c). At a daily capacity of 1% of the complete grid, the selected score captured 23.2% of total observed burned area and 21.8% of extreme events, compared with 18.7% and 12.0% for PoF only, corresponding to relative improvements of 24.0% and 81.6%. At 5% capacity, PoF only prioritisation captured 54.4% of total observed burned area, 45.9% of extreme events, and 59.8% of all burn events. The selected score captured 61.0%, 59.0%, and 57.3%, respectively, corresponding to relative changes of +12.1%, +28.7%, and −4.2%. By 10% capacity, the multiplicative score and PoF × fixed climatology showed near-identical burned area and extreme event capture, 80.4% and 81.6% compared with 80.3% and 81.6%, indicating that the additional value of the daily prediction over the climatology is concentrated at low to moderate capacities. Unlike in the Mediterranean, the selected score captured fewer burn events than PoF only prioritisation throughout the operationally relevant range of capacities; for example, 73.0% compared with 76.2% at 10% capacity.

The spatial components of the selected multiplicative score are illustrated for 25 October 2023 (Figure 2d–g), when extensive observed burned area clusters occurred across northern Australia (Figure 2d). The PoF prediction identified widespread elevated likelihood of fire occurrence across much of northern and eastern Australia (Figure 2e). FBA indicated large potential burning across much of the continent, including arid interior regions where fire occurrence was unlikely (Figure 2f), illustrating why FBA is informative only when weighted by occurrence probability. By combining these signals, the selected score ranking concentrated the highest-priority cells across northern Australia, in the region of the largest observed fire clusters (Figure 2g).

*Table 3 Australia score performance at the daily action capacity of 0.1% of the complete domain during 2020–2025, corresponding to a mean of 291 selected grid cells per day. Values show the total percentage of total observed burned area, extreme events, and number of events captured by each prioritisation score. Random selection denotes the expected capture under an unskilled score.*

| Priority Score (0.1% capacity) | Total Burned Area Capture (%) | Extreme Event Capture (%) | Number of Events Captured (%) |
|---|---|---|---|
| PoF only | 2.74 | 1.38 | 3.72 |
| PoF × climatology | 2.97 | 3.03 | 2.11 |
| Multiplicative | 3.82 | 3.82 | 3.22 |
| Random Selection | 0.1 | 0.1 | 0.1 |

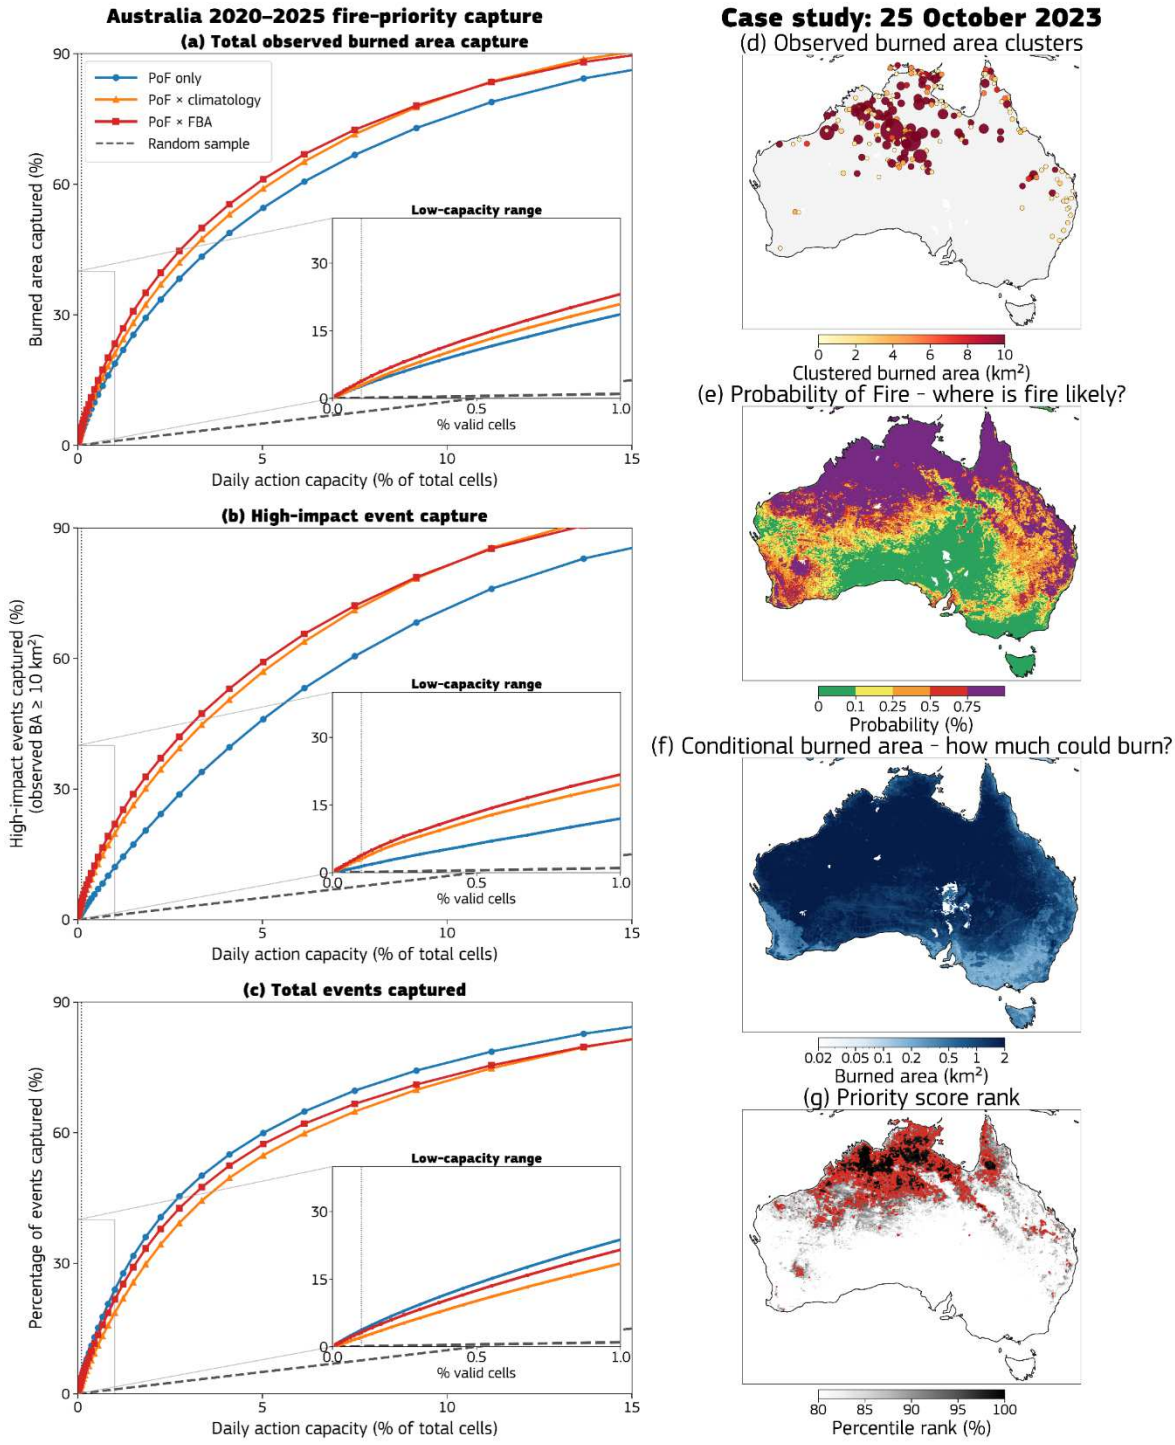



*Figure 2* ***Australia fire priority evaluation and illustrative case study.*** *(a–c) 2020–2025 capture curves for (a) total observed burned area, (b) extreme events (≥10km²/cell/day) and (c) number of gridcells with fire. Curves show the percentage captured as daily action capacity increases. Blue indicates Probability of Fire (PoF) only prioritisation; orange indicates PoF × burned area climatology from 2003-2019; red indicates the selected score, which is the multiplication of PoF with conditional burned area predictions (FBA); and dashed grey indicates random selection. Insets show the low-capacity range. (d–g) Example daily outputs for 25 October 2023 for (d) clustered observed burned area, (e) PoF predictions, (f) FBA, and (g) The multiplicative percentile-ranked priority score. In panel (g), red outlines identify contiguous regions whose scores are ≥95th percentile as an illustrative example of priority regions. The example date was selected based on the number of events, total burned area and number of extreme events.*

## Canada Prioritisation Performance

For the Canadian domain, an action capacity of 0.1% corresponds to 759 grid cells per

day. PoF only prioritisation captured 8.6% of total observed burned area, 11.5% of extreme events, and 5.6% of all burn events during 2020–2025 (Table 4). Unlike in the Mediterranean and Australia, PoF alone captured extreme events at roughly twice the rate of burn events, indicating that locations with high occurrence probability in Canada tended to coincide with the fires that went on to burn extensively. On average, approximately 245 cells burned per day against a budget of 759 selected cells. However, burning was concentrated in the fire season, and on peak days the number of burned cells exceeded the budget, so a perfect prediction could have captured at most 73.9% of burn events and 96.6% of the total burned area. The daily number of extreme events never exceeded the budget, therefore a perfect prediction could have captured all of them.

The selected Canadian formulation, PoF × FBA, captured 9.8% of total observed burned area, 14.2% of extreme events, and 5.9% of burn events. Compared with PoF only prioritisation, this represented relative improvements of 14.0% for total burned area capture, 22.9% for extreme event capture, and 4.3% for burn event capture. The relative gains in extreme event capture were smaller than in the other regions, reflecting the stronger baseline performance of PoF alone.

In contrast to both other regions, the fixed burned area climatology degraded performance relative to PoF alone. At 0.1% capacity, PoF × fixed climatology captured 3.1% of total observed burned area, 2.7% of extreme events, and 2.4% of burn events, less than half the capture achieved by PoF only for each metric. PoF × fixed climatology remained below PoF only at all capacities, and its burned area capture levelled off at approximately 46–50% between 10% and 50% capacity. This plateau suggests that roughly half of the observed burned area occurred in cells with no conditional burned area signal in the 2003–2019 climatology, which consequently received a priority score of zero. The daily FBA, which is not constrained by historical fire occurrence, avoided this limitation.

The advantage of the selected multiplicative score over PoF only prioritisation persisted across larger action capacities (Figure 3a–c). At a daily capacity of 1% of the complete grid, the selected score captured 37.1% of total observed burned area and 45.0% of extreme events, compared with 35.4% and 41.0% for PoF only, corresponding to relative improvements of 4.7% and 9.9%. At 5% capacity, PoF only prioritisation captured 68.3% of total observed burned area, 71.0% of extreme events, and 59.1% of burn events. The selected multiplicative score captured 71.7%, 79.4%, and 59.1%, respectively, corresponding to relative improvements of 5.0% and 11.9% for burned area and extreme events, with no change in burn event capture. The extreme event advantage was maintained at 10% capacity (92.4% compared with 84.4%), while burn event capture remained near-identical between the two scores across all capacities (Figure 3c).

The spatial components of selected multiplicative score are illustrated for 22 June 2023 (Figure 3d–g), during the record-breaking 2023 Canadian fire season, when the largest observed burned area clusters occurred in Quebec, with further clusters across Ontario and western Canada (Figure 3d). The PoF prediction identified elevated likelihood of fire occurrence across the western boreal region and a band extending from Ontario into

Quebec (Figure 3e), while FBA indicated relatively high potential burning across much of the forested domain (Figure 3f). By combining these signals, the multiplicative score ranking concentrated the highest-priority cells along the southern boreal band from Ontario to Quebec and across parts of western Canada, including the region of the largest observed fire clusters (Figure 3g).

*Table 4 Canada score performance at the daily action capacity of 0.1% of the complete domain during 2020–2025, corresponding to a mean of 759 selected grid cells per day. Values show the total percentage of total observed burned area, extreme events, and number of events captured by each prioritisation score. Random selection denotes the expected capture under an unskilled score.*

| Priority Score (0.1% capacity) | Total Burned Area Capture (%) | Extreme Event Capture (%) | Number of Events Captured (%) |
|---|---|---|---|
| PoF only | 8.63 | 11.52 | 5.62 |
| PoF × climatology | 3.10 | 2.74 | 2.43 |
| Multiplicative | 9.83 | 14.16 | 5.86 |
| Random Selection | 0.1 | 0.1 | 0.1 |

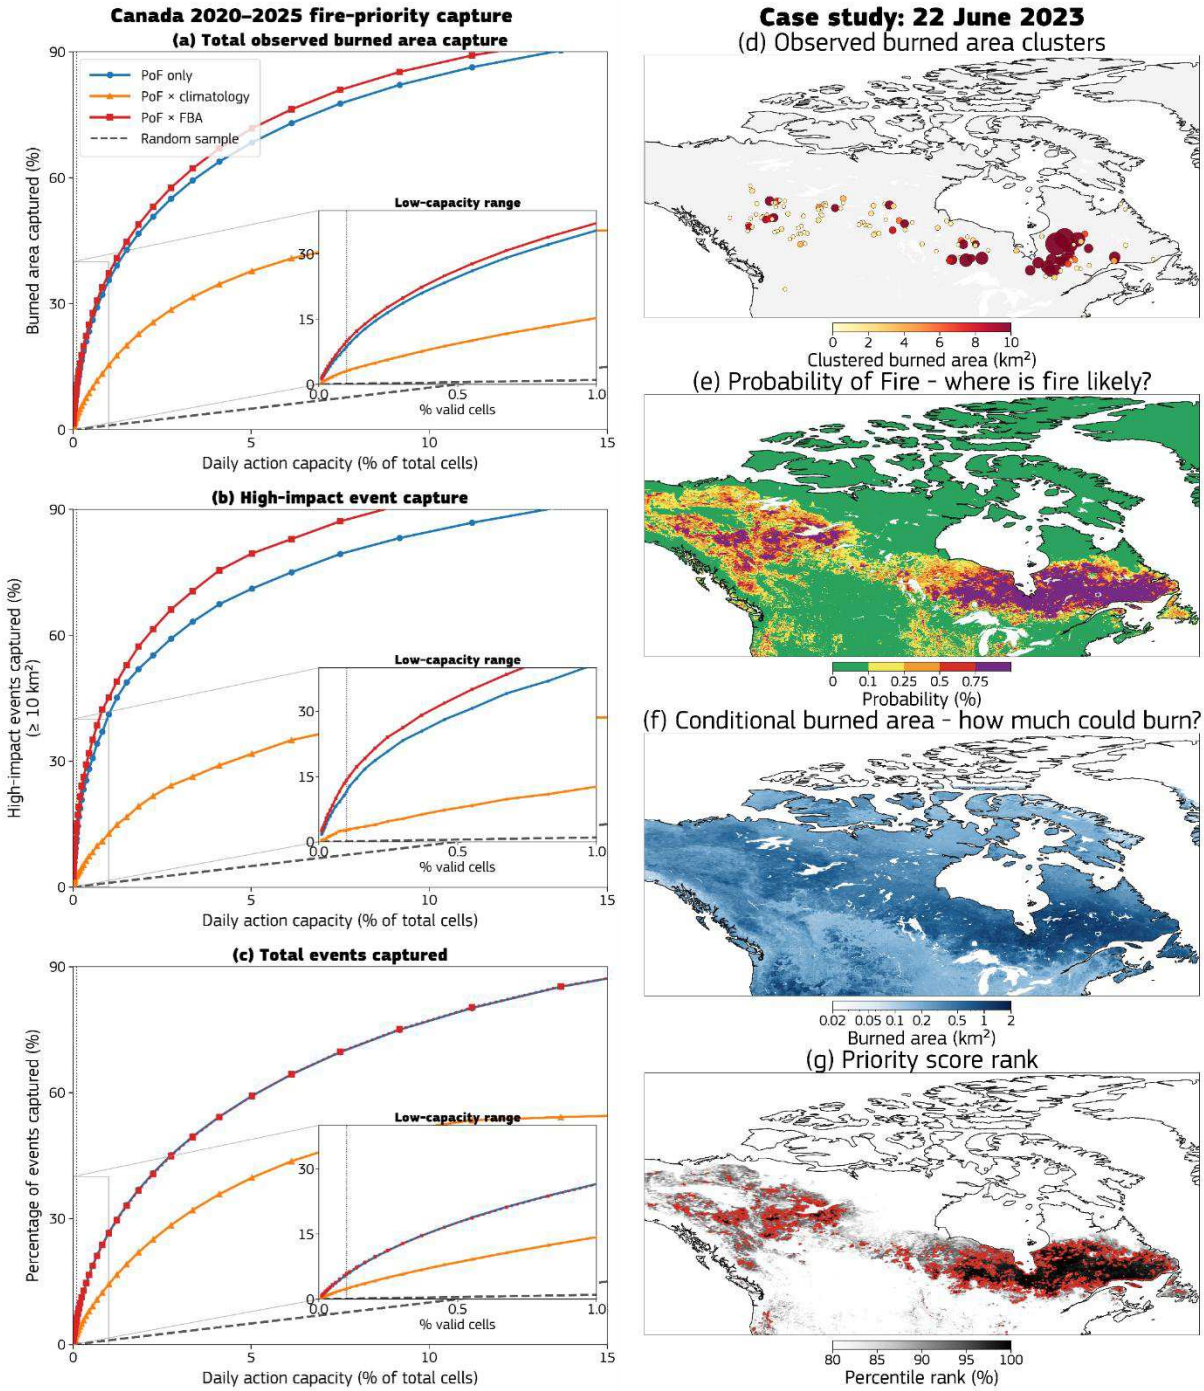



*Figure 3* ***Canada fire priority evaluation and illustrative case study.*** *(a–c) 2020–2025 capture curves for (a) total observed burned area, (b) extreme events (≥10km²/cell/day) and (c) number of gridcells with fire. Curves show the percentage captured as daily action capacity increases. Blue indicates Probability of Fire (PoF) only prioritisation; orange indicates PoF × burned area climatology from 2003-2019; red indicates the selected score, which is the multiplication of PoF with conditional burned area predictionss (FBA); and dashed grey indicates random selection. Insets show the low-capacity range. (d–g) Example daily outputs for 22 June 2023 for (d) clustered observed burned area, (e) PoF predictons, (f) FBA, and (g) The multiplicative percentile-ranked priority score. red outlines identify contiguous regions whose scores are ≥95th percentile as an illustrative example of priority regions. The example date was selected based on the number of events, total burned area and number of extreme events.*

## Cost-loss Analysis

We evaluated the net benefit of each prioritisation score relative to inaction across a continuous range of resource costs, $C_{res}$, and daily action capacities, for each domain

(Figure 4). A score yields positive net benefit where the loss reduction achieved by its selected cells outweighs the resource cost of acting on all cells up to that capacity. Because higher-ranked cells are selected first, the marginal benefit contributed by each additional selected cell declines with capacity and is greatest at the smallest capacities. Cumulative net benefit is nonetheless not maximised at the smallest capacities, as additional cells continue to add some benefit as capacity increases. This continues until less additional burned area is captured, and the marginal capture falls below the marginal cost of acting, at which point the score incurs a net cost. We examine this in three ways. First, the net benefit for each score at a single representative cost, $C_{res}$ = \$500/cell/day (Figure 4a). Second, the capacity that maximises net benefit across the full evaluated cost range (Figure 4b). Third, the value added by the selected optimum score over each comparator score as a function of capacity (Figure 4c). The perfect prediction represents an optimal burned area ranking of grid cells and forms an upper bound on every score at every cost and capacity. Notably, its optimal capacity was typically 0.2–1.0%, consistently far lower than the optimal capacity of the evaluated scores (0.2–2.8%), reflecting the strong concentration of observed large fires into a small number of grid cells once perfect knowledge of its location is available.

At $C_{res}$ = \$500/cell/day between 2020 and 2025, PoF × climatology and the selected PoF-gated score achieved the highest net benefit in the Mediterranean at action capacities of 2.25% with net benefits of \$6,645 and \$6,640 million, respectively. This benefit was considerably higher than the peak PoF only benefit of \$5,375 million occurring at capacities of 1.84% (Figure 4a). Of the three, the selected optimum scores achieved the highest net benefit of \$904 million at 0.20% capacity for Canada and \$6,640 million at 2.75% capacity for Australia. The PoF only score typically follows a similar net benefit trajectory as the selected multiplicative score across the range of action capacities with diminished absolute values for both Australia and Canada. In Canada, under this cost and loss assumption, the PoF × climatology score fails to provide a net benefit at any action capacity, consistent with its degraded performance relative to PoF only reported in Table 4.

Optimal capacity declined with increasing resource costs, $C_{res}$, across all scores and domains (Figure 4b). At the lowest evaluated cost, $C_{res}$ = \$20/cell/day, optimal capacity extended to 25.0% for PoF only in the Mediterranean and Australia, and to 17.2–23.1% for PoF × climatology and the selected optimum score across all three domains. This typically represents the point at which most or all the considerable burned area cells are captured by the score. At the highest evaluated cost, $C_{res}$ = \$10,000/cell/day, optimal capacity for every score contracted to below 0.04% in all domains. Domains differed in how much resource cost a score could sustain before ceasing to be worth deploying at any meaningful scale. In Canada, the optimal capacity of the PoF × climatology score fell below 0.01% once $C_{res}$ exceeded \$410/cell/day, compared with \$1,030/cell/day for PoF only and \$1,256/cell/day for the optimum score, indicating the selected score remained economically viable across a wider range of assumed costs than the other two scores in this domain. In Australia, all three policies sustained viability to a similar, higher cost (\$1,679–2,464/cell/day), with the multiplicative score again tolerating the highest cost before its optimal capacity collapsed.

Relative to PoF only, the added value of the multiplicative score was positive across the full evaluated capacity range in Canada ($36–557 million) and Australia ($54–2,280 million; Figure 4c). The added value from the PoF-gated score was also evident in the Mediterranean; however, it reversed at high-capacity becoming negative above approximately 19% capacity after peaking at $1,543 million near 4.3% capacity. Relative to PoF × climatology, the added value was strongly and consistently positive in Canada ($155.7–8,154.4 million), reflecting the poor performance of the climatology in this domain. In Australia it was positive across most of the range, peaking at $927.5 million near 1.9% capacity, but reversed above approximately 10.3% capacity. In the Mediterranean it oscillated narrowly between −$154.2 million and $122.5 million across several capacity thresholds, although typically an added value was provided by the PoF-gated score. Whilst no evaluated score closely approached the theoretical optimal burned area benchmark at any capacity examined, the shortfall was proportionally smallest in Canada.

Absolute net benefit differed between domains, reflecting differences in assumed loss avoided, score skill, total burned area and grid-cell count (Figure 4). Net benefit in Canada was an order of magnitude smaller than in the Mediterranean or Australia throughout. Relative to the Mediterranean, this partly reflects the lower assumed loss avoided per unit area in Canada, which was derived from insured losses and is diluted by the large fraction of burning in remote boreal forest (Supplementary Information). Relative to Australia, which has a lower assumed loss value, the difference reflects the larger number of selected cells, and therefore greater resource cost, at a given capacity in Canada, together with the smaller burned area available to offset that cost. The Australian results should be interpreted with caution, as the assumed ($C_{miss} - C_{cap}$) is a compromise between a low estimate for remote savanna burning and a substantially higher estimate for fires in populated areas. The true value added by the multiplicative score in Australia may therefore differ substantially from the figures presented here, depending on where fires are concentrated. The sensitivity of net benefit to $C_{res}$ and ($C_{miss} - C_{cap}$) across a wider range of values is shown at fixed daily action capacities in Supplementary Figure 1. In Canada and Australia, and in the Mediterranean at low and moderate capacities, the break-even line for the optimal score lay at higher resource costs, or lower loss avoided, than that for PoF × climatology. This indicates that the economic advantage over the climatology-based score did not depend on specific cost and loss values. The exception was the Mediterranean at the highest evaluated capacity (2%), where the break-even lines of the two scores nearly coincided, consistent with their similar performance at high capacity.

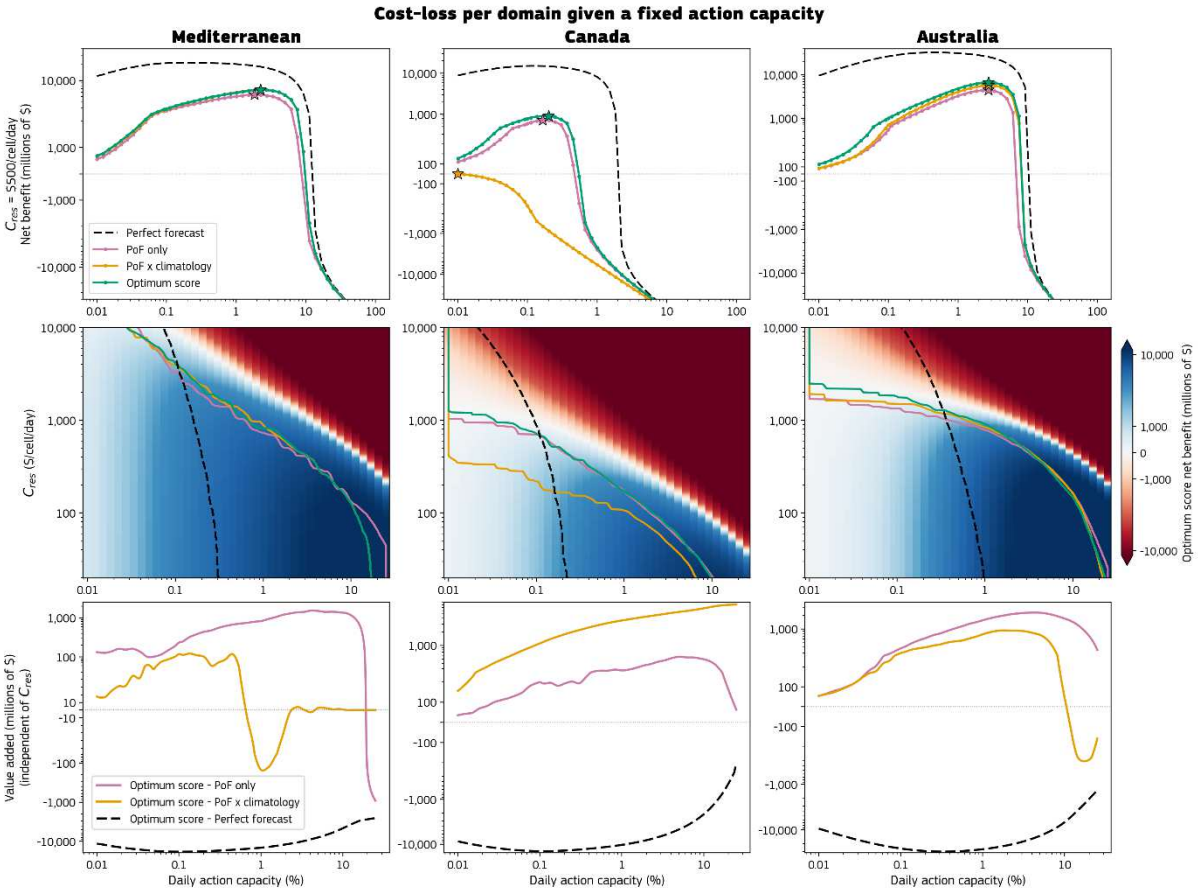



*Figure 4* ***Cost-loss value of fire priority policies across daily action capacity and resource costs.*** *Columns show the Mediterranean, Canada, and Australia. The rows show three complementary views of the same underlying quantity, net benefit relative to no action. (a) Net benefit at a single representative resource cost, $C_{res}$ = $500 per selected cell per day, as a function of daily action capacity, for multiple scores with stars marking the optimal capacity. (b) Heatmap of net benefit for the optimal score as a function of daily action capacity and resource cost. Coloured lines show the optimal capacity for each score across as a function of resource cost. (c) Value added by optimal score over each comparator score, expressed as a function of daily action capacity. Net benefit is aggregated across the full 2020–2025 evaluation period. The assumed loss avoided per km$^2$ of burned area protected, $C_{miss} - C_{cap}$, is fixed per domain.*

# Discussion

This study shows that combining daily fire occurrence predictions with information on the potential magnitude of burning improves the prioritisation of fire management and monitoring resources. Across all three regions analysed, Australia, Canada and the Mediterranean, scores that combined PoF with conditional burned area captured more observed burned area and more extreme events than PoF alone at the same daily action capacity. At a capacity of 0.1% of all domain grid cells, extreme event capture (≥10 km$^2$ day$^{-1}$) increased by 23–177% and total burned area capture by 8–40%. These gains were largest where capacity was low, which is where prioritisation decisions are hardest. This matters because demand for suppression resources increasingly arises simultaneously across regions that rely on one another for support, and sharing arrangements become strained when fire seasons overlap[15]. As fire size distribution is strongly skewed, with a small fraction of fires accounting for most burned area[16,17], occurrence-based predictions such as fire danger indices and PoF cannot on their own separate the many small fires from the few that become large. The capture framework applies the long-standing risk principle of combining likelihood with consequence[5] at daily operational

time scales. It evaluates each score by how much of the observed burning would have occurred in the locations it prioritised, given a fixed daily capacity.

The optimal score for prioritisation differed between regions, and these differences are physically interpretable. In the Mediterranean, fires tend to burn repeatedly in the same landscapes, shaped by dense human ignition sources and persistent land use[18], so historic events can inform forecasts. The fixed climatology therefore carried a strong signal, providing around half of the improvement over PoF, and FBA added value mainly by refining priorities among cells already likely to burn. In Canada, burning is dominated by large, often lightning-ignited boreal fires with long return intervals[17] and the locations that burned in 2020–2025, including during the record 2023 season[1], were poorly represented by the 2003–2019 climatology. Around half of the observed burned area fell in cells with no historical signal, and climatology-based scores performed worse than PoF alone. The daily FBA, which depends on current conditions rather than fire history, avoided this limitation. In Australia, the volume of burning far exceeded the selection budget, and FBA supplied most of the improvement. The case study nonetheless showed that FBA must be weighted by occurrence probability, because unweighted FBA indicated large potential burning across the arid and semi-arid interior, where fire occurrence was unlikely. The value of a fixed climatology therefore depends on how previous fires can inform future ones, with no single score optimal everywhere.

The score selection search also offers a methodological insight. Bounded and rank-based candidates capped FBA, which prevented a few large predictions from dominating the ranking but also removed distinction among the largest predictions. This constraint benefited the Mediterranean, where extremes are rare and raw PoF × FBA (4.0%) captured fewer than two-thirds of the extreme events captured by the PoF-gated score (6.1%). In Canada and Australia, where extremes dominate total burned area, the raw multiplicative score was preferred, suggesting that preserving the forecast tail is most valuable where extreme fires contribute noticeably to total burned area.

These gains relative to PoF came with small relative reductions in fire event capture, of 5% in the Mediterranean and 13.5% in Australia, because PoF is optimised to predict fire occurrence rather than fire magnitude. Thus, the appropriate score depends on the management objective. PoF alone may guide initial efforts on as many ignitions as possible, whereas magnitude-weighted scores guide pre-positioning against fires that could escape. This dependence on the stated objective is a general feature of resource-constrained fire management. In a conservation context, modest targets may favour allocating resources to lower-risk sites with long fire return intervals, whereas more ambitious targets may favour concentrating them where fire is most frequent[19].

The cost–loss analysis translates capture into a measure of decision value. Under assumed parameters of cost and loss, the selected prioritisation score provided the highest or near-highest net benefit in all domains and remained viable at higher resource costs than PoF alone. In Canada, for example, it remained viable up to around $1,260 per cell per day, compared with $1,030 for PoF. Both the selected score and the optimal capacity shifted with the assumed cost of action. This is consistent with evidence that the economic value of a forecast depends on the defined cost–loss structure. It also

supports the idea that conventional frameworks, which treat extremes in isolation, misrepresent the compounding losses from concurrent events[20]. Concurrent fires competing for shared resources are an example of such compounding, and the capacity constraint represents one aspect of this competition, although losses here are not modelled as compounding. The perfect-forecast benchmark peaked at consistently lower capacities than any evaluated score. This indicates that much of the attainable value lies in identifying a few high-consequence events.

The monetary values used in this study are illustrative rather than estimates of realised benefit. They depend on a single domain-wide loss value and a single cost per selected cell, which ignore spatial variation in assets, exposure and suppression cost. This is particularly relevant in Australia, where the loss value is a compromise between remote savanna burning and fires near populated areas. The analysis also treats losses as purely financial, omitting values that are difficult to monetise, such as human health, biodiversity, carbon and cultural heritage, and it assumes all burning is detrimental, although some fire, including planned and cultural burning in savanna landscapes, is ecologically beneficial. The framework is general, however, and could be parameterised with region- or agency-specific values, including spatially varying losses, non-market or conservation objectives[19], and separate costs for monitoring, pre-positioning and suppression. In that form, it could indicate both which forecast to use and how much capacity it is economically viable to commit.

Several limitations apply to the framework presented in this study, each suggesting directions for future work. The cost–loss parameters were fixed rather than treated as uncertain, and accounting for parameter uncertainty can change which allocation appears optimal. A probabilistic treatment of costs and losses would test the robustness of the preferred scores. Burned area is an imperfect proxy for loss, as similar burned areas can result in very different outcomes. Because the framework can accommodate any target representing fire consequence, exposure- or conservation-weighted targets could move the score closer to a direct measure of loss. Alternatively, variables like fire radiative power could replace burned area where suppression feasibility is the concern, since intensity thresholds limit whether direct attack is possible[21]. Observed burned area also already reflects the suppression carried out, so the analysis cannot represent how additional action would have changed fire growth. In addition, MODIS underestimates small fires[22], and the evaluation treats grid cells independently, whereas in practice resources are deployed to fires or management areas that can span multiple cells and vary in size. Finally, evaluation at longer lead times and with ensemble forecasts would show how far in advance prioritisation can usefully be informed.

In summary, adding information on the potential magnitude of burning to occurrence forecasts directs limited fire-management attention more effectively towards high impact fires. The benefit holds across contrasting fire regimes, although the best way to combine the two sources of information is region-dependent. As concurrent fire seasons increase pressure on shared resources[23], forecasts designed and evaluated explicitly for prioritisation under finite capacity offer a practical step towards the resource-planning tools required.

# Methods

## Model overview

We evaluated daily wildfire-priority scores across the Mediterranean, Canada, and Australia during 2020–2025. These regions were selected because their fire seasons repeatedly exceed domestic suppression capacity, prompting formalised international resource sharing arrangements, and because co-occurring large fires competing for shared resources are expected to become more frequent as fire seasons lengthen and overlap[15]. They also span contrasting fire regimes and management contexts, allowing us to test whether the value of combining occurrence probability with conditional burned area is region dependent. The analysis used PoF predictions[7], daily conditional burned area predictions, fixed conditional burned area climatology and observed burned area from the MODIS Collection 6.1 Burned Area product MCD64A1[22]. Domain definitions are summarised in the Supplementary Information.

PoF estimates the probability of satellite-detectable fire occurrence from daily environmental and human-related predictors. The development, predictor data, and forecast evaluation of PoF are described elsewhere[7,8].

We developed a complementary daily conditional burned area prediction (FBA) model to estimate burned area fraction at the grid-cell scale conditional on a satellite detection. Machine learning has previously been used to predict final fire size given an ignition, showing environmental conditions contain information on eventual fire size[24]. Here, we extend this conditional approach from individual fires to daily, gridded predictions across continental domains. The model target was burned area fraction derived from the MODIS Collection 6.1 Burned Area product MCD64A1[22]. Training data covered 2003–2019 omitting all gridcell-days without burn events. The model therefore represented the potential magnitude of burning of all gridcells conditional on fire occurrence, rather than a discrete representation of whether a fire occurs.

The conditional burned-area model was fitted using an extreme gradient boosting regression model, XGBoost[25], with a Tweedie objective function (variance power = 1.5), suitable for the strongly right-skewed, non-negative distribution of burned area fractions. The model used 1,200 trees, a learning rate of 0.02, maximum tree depth of 6, minimum child weight of 60, L2 regularisation of 15, L1 regularisation of 0.5, and histogram-based tree construction. Low- and high-vegetation land-cover classes were represented as categorical predictors. The hyperparameter tuning was performed based on a range of viable options, with results indicating that the configuration described above yielded the most suitable tail to capture extreme events over the training period.

Predictors characterised daily weather, fuel moisture, vegetation and fuel conditions, topography, lightning, land cover, and human influences. Daily meteorological predictors included precipitation, 2m air temperature, 2m dewpoint temperature, and wind speed taken from ERA5-Land[26], and lightning occurrence[27]. Fuel and vegetation predictors included dead-fuel-moisture indicators, live-fuel moisture, fuel-load estimates, and monthly low- and high-vegetation leaf-area index[28]. Static or annually varying predictors included low- and high-vegetation cover, crop cover, slope, road

density, population density, urbanisation, night-time lights, gross domestic product, governance, human-development indicators, and related socioeconomic variables[3,29]. All variables and observed burned area data are regridded to a global ~9 km resolution horizontal resolution.

To reduce computational demands while retaining seasonal and interannual representation, a random 50% sample of all non-zero grid-cell-days were retained within each month before aggregating the 2003–2019 training dataset. Training observations above the 90th percentile and up to the 99.9th percentile of the retained cells were assigned a weight of 0.3, whereas observations above the 99.9th percentile were assigned a weight of 15. Observations at or below the 90th percentile retained a weight of 1. This weighting increased emphasis on the rarest, largest burned-area events while limiting the influence of the more numerous intermediate upper-tail observations thus preventing a general overestimation.

Year-specific observed burned area was used only for evaluation and was not used to construct priority scores. Fixed conditional burned area climatology was calculated from the 2003–2019 period as the mean burned area conditional on observed burning in each grid cell. Missing or non-finite climatological values were assigned zero, representing no historical conditional fire events recorded. Climatological burned area and FBA were combined with PoF in multiple ways to produce the candidate priority scores.

## Priority-score construction and selection

Each day, priority scores ranked all domain grid cells according to their relative priority for fire management attention. At a fixed daily action capacity, the highest ranked cells were selected first. Selection represents possible targeting of enhanced monitoring, preparedness, resource pre-positioning, or situational awareness; it does not represent a literal deployment instruction or guarantee that suppression action would occur successfully.

The reference score ranked cells using daily probability of fire occurrence (PoF). We tested whether this ranking could be improved by incorporating fixed conditional burned-area climatology and daily conditional burned-area predictions (FBA). Candidate scores included direct multiplicative PoF–burned-area scores, bounded climatology-based residual adjustments, PoF-gated prediction adjustments, two-stage PoF-screening scores, daily percentile-rank blends, geometric blends, and PoF-weighted climatology–prediction mixtures. The full candidate inventory and parameter ranges are reported in the Supplementary Information.

One direct multiplicative candidate combined daily fire-occurrence probability and the raw conditional burned area prediction:

$$S_i = PoF \times FBA$$

Where $S_i$ is the priority score for grid cell $i$, $PoF_i$ is daily probability of fire and $FBA$ is the conditional burned area prediction. This candidate used the raw prediction without a PoF

gate, climatological burned-area term, capping, or transformation.

We also evaluated a PoF-gated climatology–prediction candidate that combined PoF and fixed conditional burned-area climatology, with a bounded FBA adjustment applied only among cells with the highest daily PoF:

$$S_i = PoF_i \times BA_{climatology,i} \times M_i$$

Where $BA_{climatology,i}$ is fixed conditional burned area climatology based on past observations, and $M_i$ is a multiplier informed by $FBA$. The multiplier was defined as:

$$M_i = 1 + \lambda\left(2 \times FBA_{norm,i} - 1\right), \qquad i \in G\,,\ 1\,, i \notin G\,,$$

where $G$ denotes a daily PoF gate, $\lambda$ controls the strength of the FBA adjustment, and $FBA_{norm,i}$ is the normalised FBA, bounded between 0 and 1. This formulation allows FBA to refine the relative priority of locations already identified as having high fire likelihood, while retaining the underlying PoF × climatology score outside the gate.

For candidate scores requiring a bounded relative FBA signal, finite positive daily FBA values were capped at the daily 80th percentile, transformed using a fourth-root transformation, divided by the transformed 80th-percentile value, and constrained to the interval [0,1]. This reduced the influence of a small number of extreme predictions on priority rankings.

Candidate scores were evaluated separately within each study region at a primary daily action capacity corresponding to 0.1% of the complete study-domain grid. On each day, this number of cells was selected from the entire domain. Within each region, candidates were ranked using the equally weighted mean of pooled total burned-area capture and pooled high-burned-area cell-day capture. PoF-gated candidates used gates corresponding to the highest 1%, 2%, 5%, 10%, and 20% of daily PoF values; bounded-adjustment strengths included λ = 0.05 and 0.10–1.00 in increments of 0.10. The highest-ranked configuration in each region was retained for subsequent capacity-curve evaluation.

## Priority-score evaluation

Priority scores selected a fixed number of valid cells each day, corresponding to 0.1% of the complete study-domain grid. This represents a low-capacity setting in which operational attention can be directed to only a small proportion of the landscape. We also assessed performance across a broader range of daily action capacities to characterise how score differences changed as capacity increased.

We evaluated three metrics. Total burned area capture was defined as the observed burned area within selected cells divided by total observed burned area across all forecast-valid cells. Extreme event capture was defined as the proportion of cell-days with observed burned area ≥10 $km^2$ within the selected cells. Finally, burned-cell capture was defined as the proportion of cell-days with any observed burning. This final metric

was retained as an additional diagnostic because identifying the occurrence of burned cells aligns with the PoF objective.

## Cost–loss analysis

Classical cost–loss frameworks quantify the value of a prediction for a single, repeated binary decision of whether to take a fixed-cost protective action against a fixed loss if an event occurs and no action is taken[30,31]. We extend this framework to the capacity-constrained, multi-location decision problem addressed here, in which a fixed number of grid cells are selected for attention each day. This connects our prioritisation scores to an operationally interpretable measure of decision value and follows recent applications of cost–loss relating to complex, spatially distributed hazards[20,32].

We define three quantities. Each selected grid cell incurs a fixed cost, $C_{res}$, representing the cost of directing monitoring, preparedness, or suppression resources to that cell for one day, irrespective of whether a fire subsequently occurs there. Each cell that burns without having been selected incurs a full loss, $C_{miss}$, per km$^2$ of burned area. Each cell that burns and was selected incurs a reduced loss, $C_{cap}$, per km$^2$ of burned area, reflecting the effect of prior positioning or attention on limiting eventual fire size or impact, with $C_{miss} > C_{cap}$ by construction. The quantity $(C_{miss} - C_{cap})$ therefore represents the loss avoided per km$^2$ of burned area in a cell that was correctly selected.

For each score, a constant number of grid cells is selected each day throughout the 2020–2025 evaluation period, based on a fixed daily action capacity. The cumulative net benefit of that score relative to taking no action, aggregated across the whole evaluation period, is:

$$Net(N) = \left[\sum_{i=1}^{N} (C_{miss} - C_{cap}) \times BA_i\right] - C_{res} \times N$$

where $BA_i$ is the observed burned area in the i-th selected cell-day and $N$ is the total number of cell-day selections pooled across the 2020–2025 evaluation period at a given action capacity.

Because higher-ranked cells are selected first, the cumulative burned area captured should result in diminishing returns as a function of N for every score, with successive cells adding progressively less burned area capture. An optimum action capacity will occur for each score for each configuration of, $C_{res}$, $C_{miss}$, and $C_{cap}$, at which a given score selection yields the greatest net benefit relative to inaction. We evaluated $Net(N)$ across a range of $C_{res}$ values, which represent scenarios from low-cost monitoring actions to higher-cost suppression responses, for each domain. This is done to identify the optimum daily action capacities, alongside a theoretical prediction with perfect knowledge of where burning would occur.

Because published cost data specific to grid-cell-level fire monitoring and pre-positioning are not available, $(C_{miss} - C_{cap})$ was fixed per domain using generalised estimates of wildfire suppression and damage costs per unit area (see Supplementary

Information), and $C_{res}$ was varied across a broad range spanning the operating costs of relatively cheap monitoring to more costly suppression assets. This analysis is illustrative rather than a literal cost–benefit accounting of any specific fire-management action. In practice, monitoring and suppression differ substantially in their operational mechanisms, response timescales, and costs, and neither is fully captured by the single parameter $C_{res}$ used here. The purpose of this analysis is not to recommend a specific operational cost threshold, but to demonstrate how the prioritisation scores developed in this study could, in principle, be translated into an economically interpretable measure of operational value, and how that value depends on the relative cost of action and the magnitude of loss it is assumed to avoid.

## Funding

This work was funded by the Copernicus Emergency Management Service contract no. 942604 between the Joint Research Centre and ECMWF. JM is funded through The CoFlame European Union project, grant agreement no. 101296262.

## References

1. Jones, M. W. *et al.* State of Wildfires 2023–2024. *Earth System Science Data* **16**, 3601–3685 (2024).
2. Kelley, D. I. *et al.* State of Wildfires 2024–2025. *Earth System Science Data* **17**, 5377–5488 (2025).
3. Burton, C. *et al.* State of Wildfires 2025–2026. Preprint at https://doi.org/10.5194/essd-2026-587 (2026).
4. Jain, P. *et al.* Drivers and Impacts of the Record-Breaking 2023 Wildfire Season in Canada. *Nat Commun* **15**, 6764 (2024).
5. Thompson, M. P., Scott, J., Helmbrecht, D. & Calkin, D. E. Integrated wildfire risk assessment: Framework development and application on the Lewis and Clark National Forest in Montana, USA. *Integrated Environmental Assessment and Management* **9**, 329–342 (2013).
6. Ntaimo, L., Arrubla, J. A. G., Stripling, C., Young, J. & Spencer, T. A stochastic programming standard response model for wildfire initial attack planning. *Can. J. For. Res.* **42**, 987–1001 (2012).
7. McNorton, J. R., Di Giuseppe, F., Pinnington, E., Chantry, M. & Barnard, C. A Global Probability-Of-Fire (PoF) Forecast. *Geophysical Research Letters* **51**, e2023GL107929 (2024).
8. Di Giuseppe, F., McNorton, J., Lombardi, A. & Wetterhall, F. Global data-driven prediction of fire

activity. *Nat Commun* **16**, 2918 (2025).

9. Bowman, D. M. J. S. *et al.* Fire in the Earth System. *Science* **324**, 481–484 (2009).
10. Johnston, F. H. *et al.* Unprecedented health costs of smoke-related PM2.5 from the 2019–20 Australian megafires. *Nat Sustain* **4**, 42–47 (2020).
11. Van Der Werf, G. R. *et al.* Global fire emissions estimates during 1997–2016. *Earth Syst. Sci. Data* **9**, 697–720 (2017).
12. Tedim, F. *et al.* Defining Extreme Wildfire Events: Difficulties, Challenges, and Impacts. *Fire* **1**, 9 (2018).
13. Krawchuk, M. A. & Moritz, M. A. Constraints on global fire activity vary across a resource gradient. *Ecology* **92**, 121–132 (2011).
14. Scott, J. H., Thompson, M. P. & Calkin, D. E. *A Wildfire Risk Assessment Framework for Land and Resource Management*. RMRS-GTR-315 https://research.fs.usda.gov/treesearch/56265 (2013) doi:10.2737/RMRS-GTR-315.
15. Bloem, S., Cullen, A. C., Mearns, L. O. & Abatzoglou, J. T. The Role of International Resource Sharing Arrangements in Managing Fire in the Face of Climate Change. *Fire* **5**, 88 (2022).
16. Strauss, D., Bednar, L. & Mees, R. Do One Percent of the Forest Fires Cause Ninety-Nine Percent of the Damage? *Forest Science* **35**, 319–328 (1989).
17. Stocks, B. J. *et al.* Large forest fires in Canada, 1959–1997. *Journal of Geophysical Research: Atmospheres* **107**, FFR 5-12 (2002).
18. Vilar, L., Camia, A., San-Miguel-Ayanz, J. & Martín, M. P. Modeling temporal changes in human-caused wildfires in Mediterranean Europe based on Land Use-Land Cover interfaces. *Forest Ecology and Management* **378**, 68–78 (2016).
19. Lawton, P. J. *et al.* Robust allocation of fire management resources towards threatened species conservation under uncertainty. *Ecological Modelling* **517**, 111582 (2026).
20. Olivetti, L., Messori, G., Avner, P. & Hallegatte, S. Assessing the real-world economic value of weather forecasts for compounding extremes. *Environ. Res. Lett.* **21**, 154048 (2026).

21. Wotton, B. M., Flannigan, M. D. & Marshall, G. A. Potential climate change impacts on fire intensity and key wildfire suppression thresholds in Canada. *Environ. Res. Lett.* **12**, 095003 (2017).
22. Giglio, L., Boschetti, L., Roy, D. P., Humber, M. L. & Justice, C. O. The Collection 6 MODIS burned area mapping algorithm and product. *Remote Sensing of Environment* **217**, 72–85 (2018).
23. Abatzoglou, J. T., Juang, C. S., Williams, A. P., Kolden, C. A. & Westerling, A. L. Increasing Synchronous Fire Danger in Forests of the Western United States. *Geophysical Research Letters* **48**, e2020GL091377 (2021).
24. Coffield, S. R. *et al.* Machine learning to predict final fire size at the time of ignition. *Int J Wildland Fire* **28**, 861–873 (2019).
25. Chen, T. & Guestrin, C. XGBoost: A Scalable Tree Boosting System. in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 785–794 (Association for Computing Machinery, New York, NY, USA, 2016). doi:10.1145/2939672.2939785.
26. Muñoz-Sabater, J. *et al.* ERA5-Land: a state-of-the-art global reanalysis dataset for land applications. *Earth Syst. Sci. Data* **13**, 4349–4383 (2021).
27. Lopez, P. A Lightning Parameterization for the ECMWF Integrated Forecasting System. *Monthly Weather Review* **144**, 3057–3075 (2016).
28. McNorton, J. R. & Di Giuseppe, F. A global fuel characteristic model and dataset for wildfire prediction. *Biogeosciences* **21**, 279–300 (2024).
29. McNorton, J., Comyn-Platt, E., Di Giuseppe, F., Cagnazzo, C. & Barnard, C. Global fuel and burned area reconstructions and projections for 1950–2100. Preprint at https://doi.org/10.5194/egusphere-2026-5374 (2026).
30. Murphy, A. H. The Value of Climatological, Categorical and Probabilistic Forecasts in the Cost-Loss Ratio Situation. *Mon. Wea. Rev.* **105**, 803–816 (1977).
31. Richardson, D. S. Skill and relative economic value of the ECMWF ensemble prediction system.

805 *Quart J Royal Meteoro Soc* **126**, 649–667 (2000).
806 32. Busker, T., De Moel, H., Van Den Hurk, B. & Aerts, J. C. J. H. Impact-based seasonal rainfall
807 forecasting to trigger early action for droughts. *Science of The Total Environment* **898**, 165506
808 (2023).
809 33. European Commission. Forest fires (DG ECHO factsheet / EU wildfire prevention). https://civil-
810 protection-humanitarian-aid.ec.europa.eu/what/civil-protection/forest-fires_en (2023).
811 34. San-Miguel-Ayanz, J. *et al. Forest Fires in Europe, Middle East and North Africa 2023*.
812 https://effis.jrc.ec.europa.eu/reports-and-publications/annual-fire-reports (2024).
813 35. Hanes, C. C. *et al.* Fire-regime changes in Canada over the last half century. *Canadian Journal of*
814 *Forest Research* **49**, 256–269 (2019).
815 36. Natural Resources Canada. Canadian National Fire Database (CNFDB).
816 https://cwfis.cfs.nrcan.gc.ca/ha/nfdb (2026).
817 37. Insurance Bureau of Canada. Facts of the Property and Casualty Insurance Industry in Canada.
818 https://www.ibc.ca/insurance-basics/facts-of-the-industry (2026).
819 38. Insurance Bureau of Canada. Northern Alberta wildfire costliest insured natural disaster in
820 Canadian history. https://www.ibc.ca (2016).
821 39. Russell-Smith, J. *et al.* Bushfires 'down under': patterns and implications of contemporary
822 Australian landscape burning. *International Journal of Wildland Fire* **16**, 361–377 (2007).
823 40. Commonwealth of Australia. *Royal Commission into National Natural Disaster Arrangements:*
824 *Report*. https://www.royalcommission.gov.au/natural-disasters/report (2020).
825 41. Boer, M. M., Resco de Dios, V. & Bradstock, R. A. Unprecedented burn area of Australian mega
826 forest fires. *Nature Climate Change* **10**, 171–172 (2020).
827 42. Insurance Council of Australia. Catastrophe data: 2019–20 bushfire season insured losses.
828 https://insurancecouncil.com.au/industry-members/data-hub/ (2020).
829
830

# Supplementary Information

## Domains

The three regional domains were bounded in rectangles used for evaluation. These were defined on a common 0.07° (~9 km) global grid.

*Supplementary Table 1.* ***Region selection summary****. Bounded regions for the three domains that were selected for score prioritisation cost-loss analysis and the number of valid grid cells for each domain.*

| Domain | Latitude range | Longitude range | Grid cells (per day) |
|---|---|---|---|
| Mediterranean | 30°N–46°N | 10°W–40°E | ~161,851 |
| Canada | 41.5°N–83.5°N | 141.4°W–52.5°W | ~758,862 |
| Australia | 44°S–10°S | 112°E–154°E | ~290,400 |

## Prioritisation Scores

Candidate scores were constructed from PoF, FBA and the fixed burned-area climatology ($BA_{clim}$). Candidate score families were built around multiple selection methods, each calculated across the eligible cells of a domain.

One such method was a bounded FBA, used by several candidate families to prevent a few extreme forecasts from dominating the ranking. FBA was capped at q80, the daily 80th percentile of FBA values, transformed with a fourth root and divided by the transformed cap, giving a normalised quantity:

$$FBA_{norm} = \frac{[min(FBA, q80)]^{0.25}}{q80^{0.25}}$$

The second was the daily percentile rank, R, used by candidates that combine inputs through their ranks rather than their values, so that differences in units do not affect the score. R(x) is the percentile rank of x among eligible cells on that day, from 0 (lowest) to 1 (highest), with tied cells assigned their average rank. Ranks were calculated for PoF, $BA_{clim}$ and $FBA_{norm}$.

The third was a PoF gate, $G_f$, used by candidates in which FBA adjusts the ranking only where fire is already likely. $G_f$ is the set of eligible cells within the top fraction f of daily PoF values; for example, $G_{f,1\%}$ contains the 1% of cells with the highest PoF that day. The remaining selection methods are summarised in Supplementary table 2.

Each day, every candidate score selected its highest-scoring eligible cells up to a daily action capacity of 0.1% of the complete domain grid, which equated to 162, 759 and 291 cells in the Mediterranean, Canada and Australia, respectively (Supplementary Table 1). In total, 197 candidate scores were tested in each domain (Supplementary Table 2).

*Supplementary Table 2.* ***Candidate priority-score families.*** *All the candidate scores evaluated with their parameter values. λ values tested: 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0.*

| Family | Score | Parameters tested | Score Count |
|---|---|---|---|
| Direct | PoF, PoF × $BA_{clim}$, PoF × FBA, PoF × $FBA_{norm}$ | none | 4 |
| Bounded residual | PoF × $BA_{clim}$ × [1 + λ(2Z − 1)], with Z = $FBA_{norm}$ or Z = R($FBA_{norm}$) | All $\lambda$ for each Z | 22 |
| PoF-gated FBA adjustment | PoF × $BA_{clim}$ × [1 + λ(2$FBA_{norm}$ − 1)] inside $G_f$; PoF × $BA_{clim}$ outside | f: 1, 2, 5, 10, 20%; All $\lambda$ | 55 |
| Two-stage screening | Within $G_f$, rank by $BA_{clim}$, or $BA_{clim}$ × [1 + λ(2$FBA_{norm}$ − 1)] | f: 0.2, 0.5, 1, 2, 5, 10%; All $\lambda$ | 72 |
| Percentile-rank blend | $W_p$ R(PoF) + $W_c$ R($BA_{clim}$) + $W_m$ R($FBA_{norm}$) | $W_{p/c/m}$: 0.8/0.15/0.05, 0.7/0.2/0.1, 0.6/0.3/0.1, 0.6/0.2/0.2, 0.5/0.4/0.1, 0.5/0.3/0.2, 0.4/0.4/0.2, 0.4/0.3/0.3 | 8 |
| Geometric blend | $PoF^A$ × R($BA_{clim}$)$^B$ × $FBA_{norm}^C$, each term floored at $10^{-6}$ | A/B/C: 1/1/0.025, 1/1/0.05, 1/1/0.1, 1/1/0.2, 1/0.75/0.05, 1/0.75/0.1, 1/1.25/0.05, 1/1.25/0.1, 1.25/1/0.05, 1.25/1/0.1, 0.75/1/0.05, 0.75/1/0.1 | 12 |
| PoF with weighted FBA | (i) PoF × $M_α$; (ii) $M_α$ ranked within $G_f$, where $M_α$ = (1 − α) × R($BA_{clim}$) + α R($FBA_{norm}$) | α: 0.05, 0.1, 0.2, 0.3, 0.4, 0.5; f: 1, 5, 10% | 24 |

The candidate families differ in how strongly FBA can reorder the ranking. In the direct family, FBA acts at full magnitude, whereas in the gated and screened families it reorders only cells that PoF already identifies as relatively likely to experience wildfires, limiting the influence of large FBA values where fire is improbable. Because FBA is capped at q80, candidates using $FBA_{norm}$ or its rank do not distinguish among the largest 20% of daily forecasts. In addition, because tied cells share an average rank, cells with zero climatology receive a non-zero R($BA_{clim}$), unlike for PoF × $BA_{clim}$, the rank-based families therefore do not exclude locations with no recorded fire history.

Each candidate was evaluated at 0.1% daily capacity using the 2020–2025 period and ranked by the equally weighted mean of total burned area capture and extreme event capture (≥10 $km^2\,day^{-1}$). The highest ranked candidate in each domain was retained for further analysis (Table 1). Because candidates were searched and evaluated over the same period, the capture rates of the selected scores should be regarded as optimistic development estimates. PoF × FBA, selected in Canada and Australia, has no tuned parameters, whereas the PoF-gated score for the Mediterranean has two (f = 1%, λ = 0.8).

The selected scores were compared with four reference scores: PoF only, PoF × $BA_{clim}$, random selection, in which the same number of eligible cells was drawn uniformly each day to provide a no-skill floor, and a perfect forecast, in which cells were ranked by observed burned area to provide an upper bound.

## Burned Area Value

The true cost of a wildfire is specific to each event. The 2025 Los Angeles fires, for example, were small in area compared with extreme boreal fires but considerably more

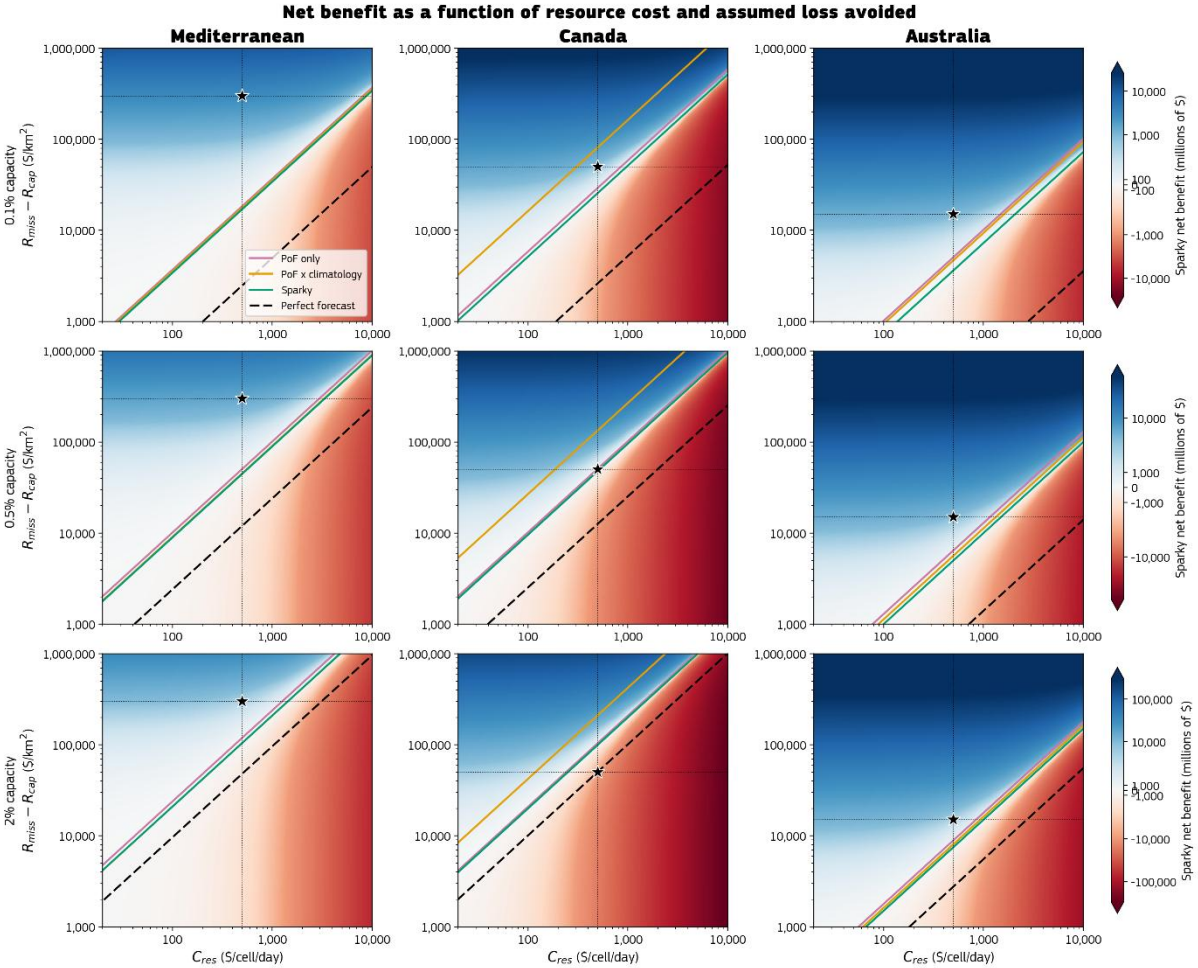
costly. The values below should therefore be read as illustrative domain averages, not as the loss per unit area of any specific fire. $C_{miss} - C_{cap}$ represents the economic loss avoided by limiting burning, not the cost of the suppression effort itself. It was set per domain from published estimates of wildfire losses and burned area, converted to US dollars at approximate exchange rates as of September 2026.

For the Mediterranean, the European Commission estimates that wildfires cause around €2 billion in economic losses per year in the EU[33]. Over the 0.5–1.0 million ha burned annually across the EU and neighbouring EFFIS-monitored countries[34], this gives approximately \$220,000–460,000 per km$^2$. We adopted \$300,000 per km$^2$ as a representative value.

For Canada, the national average area burned is approximately 2–2.5 million ha per year[35,36]. Average annual insured wildfire losses of the order of \$500 million[37] equating to a national estimate of approximately \$20,000–25,000 per km$^2$. This average is diluted by the large fraction of burning in remote boreal forest with little insured value. By contrast, the 2016 Fort McMurray fire caused insured losses of approximately \$2.6 billion[38] (\$440,000 per km$^2$). This is likely an overestimate for typical high-value fires, since nearly all financial loss was concentrated in the townsite, a small fraction of the fire footprint. We adopted \$50,000 per km$^2$, weighted towards the national blended estimate while allowing for the substantially higher losses possible near populated areas.

For Australia, burned area is dominated by frequent, extensive fires in the remote northern savannas[39], so national financial losses divided by national burned area would give a very low per-km$^2$ value. Financial losses are instead concentrated in the populated south-east. The 2019–20 Black Summer fires, which burned more than 24 million ha nationally[40], including around 5.8 million ha of temperate forest in the south-east[41] , generated insured losses of approximately AUD 2.3 billion[42]. Expressed per unit area, this ranges from ~\$6,000 per km$^2$ over the full national burned area to \$15,000–25,000 per km$^2$ over the south-eastern fire extent. We adopted \$15,000 per km$^2$ as a compromise. It lies at the lower end of the south-eastern range, reflecting that most burned area in Australia is remote and of low relative financial value. Insured losses exclude uninsured, health, ecological and economy-wide costs, so this value is conservative.

Across all three domains, we treat these values and $C_{res}$ as highly uncertain parameters, and we recommend that they be revisited with region-specific cost data where available.




*Supplementary Figure 1.* ***Sensitivity of net benefit to assumed resource cost and loss avoided, at multiple daily action capacities.*** *Columns show the Mediterranean, Canada, and Australia domain and rows show three fixed daily action capacities, 0.1%, 0.5%, and 2%, of the complete domain grid. Capacity is held constant within each panel, and net benefit is evaluated across a continuous range of resource cost, $C_{res}$ (x-axis), and assumed loss avoided per km² of burned area protected (y-axis). Color indicates the net benefit for the selected score and coloured lines show the break-even combination of $C_{res}$ and $C_{miss} - C_{cap}$ at which net benefit equals zero for each score, with the perfect forecast (dashed) forming an upper bound. The star and dotted crosshairs indicate the assumed values used in the study. Net benefit is aggregated across the full 2020–2025 evaluation period.*

## Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformation.docx