Authors' response to reviewer's comments

Portable weather radar for wildfire intelligence: observations, diagnostics, and a roadmap toward operations

We would like to thank the editor and the two reviewers for their very constructive comments on our manuscript. We received genuine insights, which have significantly contributed to increasing the manuscript quality and potential impact. To improve the clarity in our responses we have numbered the reviewers' comments: for example, the comment 1 from reviewer 1 is listed as R1C1 and will refer to these comments as such in the following.

Reviewer 1

This is a timely and innovative manuscript addressing an important observational gap in wildfire science and operational fire intelligence. The paper makes a strong case for portable weather radar as a complementary observing platform capable of resolving plume structure, pyroconvective dynamics, pyrometeor transport, and low-level wind changes at spatial and temporal scales that are generally inaccessible to satellites, numerical models, and sparse surface observations. The integration of a literature synthesis, radar-derived diagnostics, a recent demonstration case, and a Technology Readiness Level framework is particularly valuable. The manuscript is scientifically interesting and contains several impressive observational examples. I particularly appreciated the effort to connect fundamental radar science with practical deployment constraints, fire-agency workflows, multisensor integration, and coupled fire–atmosphere modelling. The Wallangarra case study and the proposed operational roadmap provide a useful foundation for future research and pre-operational development.

I am positive about the manuscript and support publication following revision. The main improvements required concern the overall structure, methodological transparency, consistency between claims and evidence, clearer separation between review material and original analysis, and substantial language and presentation editing. At present, some sections read as a synthesis paper, others as a methods or case-study paper, and others as a perspective or operational concept article. Clarifying this identity would significantly strengthen the contribution.

Major comments:

- The manuscript combines four substantial components: i) a review of approximately two decades of portable radar observations, ii)a synthesis of radar-derived wildfire diagnostics, iii)original analysis of the Wallangarra case, and iv) a TRL assessment and operational roadmap. All four components are valuable, but their relative roles are not sufficiently clear. The title and abstract emphasize portable radar for wildfire intelligence, whereas large parts of the paper discuss operational fixed-site radar examples, previously published cases, model evaluation, and conceptual operational workflows. I suggest that the authors explicitly define the manuscript as either i) a review and perspective paper supported by a demonstration case, ii) or an original research paper presenting a diagnostic framework, with the review serving as context.

R1C1: We thank the reviewer for this important comment. We agree that the previous version did not make the relative roles of the review, diagnostic synthesis, Wallangarra analysis, and operational roadmap sufficiently explicit.

We have therefore reframed the manuscript as a review and perspective study supported by a demonstration case. This positioning is now stated clearly in the title, abstract, and at the end of the Introduction. The revised manuscript distinguishes the four components as follows: the literature review establishes the observational evidence base from approximately two decades of portable-radar studies; the diagnostic synthesis organises reported radar signatures into a wildfire-intelligence framework; the Wallangarra case illustrates the application of selected diagnostics during a recent mobile-radar deployment; and the component-specific TRL assessment and roadmap identify the requirements for future operational integration.

We have also revised the relevant section headings and introductory text to clarify that previously published and fixed-site radar examples are included to support the broader synthesis, whereas the Wallangarra analysis is presented as an illustrative demonstration rather than as a comprehensive validation of the full framework.

- Several examples are taken from earlier publications, while the Wallangarra analysis appears to contain the main new results. However, the distinction is sometimes difficult to follow. The authors should clearly identify for every figure and analysis whether it is reproduced or adapted from a previous publication, newly processed from an existing dataset, newly analysed for this manuscript, or presented only as a conceptual illustration.

R1C2: Thank you for this important comment. We agree that the provenance of the figures and analyses was not always sufficiently explicit, particularly because the manuscript combines previously published examples, reprocessing of existing datasets, new analysis undertaken for this study, and conceptual material. We have therefore reviewed the manuscript systematically and now identify the status of each figure and analysis in the relevant caption and, where appropriate, in the main text.

Specifically, each item is now described as one of the following:
(i) reproduced from a previous publication;
(ii) adapted from a previous publication;
(iii) newly processed from an existing dataset;
(iv) newly analysed for this manuscript; or
(v) presented as a conceptual illustration.

We have also clarified that the Wallangarra case constitutes the principal new demonstration analysis in the manuscript, while the other published cases are included primarily to support the review and synthesis of portable-radar applications for wildfire intelligence. These revisions should make the respective roles of the review material, demonstration case, and conceptual operational framework more transparent.

A statement identifying the origin and analytical status of each figure has also been added directly to the figure captions so that this information is available without requiring the reader to infer it from the surrounding text.

- The Wallangarra case is central to the manuscript, but much of the methodological information is deferred to the Supplement. The main text should include sufficient detail to understand and evaluate the analysis independently. For example, how plume boundaries were identified, how the plume origin and principal axis were calculated, the radar volume and temporal sampling, quality-control and attenuation-correction procedures, treatment of ground clutter and precipitation contamination, uncertainty associated with the calculated plume volume and various metrics.

R1C3: Thank you for this important comment. We agree that the main manuscript previously relied too heavily on the Supplement for the Wallangarra methodology, despite this case providing the principal new demonstration analysis. We have therefore added a dedicated methodological subsection at the beginning of Section 5. This subsection now describes the radar sampling strategy, including the 2-min full-volume interval and 12 elevation scans; the principal quality-control and spatial restrictions; the two-stage Gaussian Mixture Model classification used to distinguish ground clutter, precipitation, pyrometeors, and Enhanced Pyrometeor Clusters; the definition of the three-dimensional plume mask and boundary; the Cartesian gridding and calculation of plume volume; and the derivation of plume orientation, origin, principal axis, axis ratio, and EPC metrics.

We have also clarified the treatment of near-radar artefacts, ground clutter, and precipitation contamination, and now state explicitly whether and how X-band attenuation correction was applied. The revised text distinguishes plume-classified radar echo volume from the complete physical smoke-plume volume and discusses uncertainties associated with radar sensitivity, beam broadening, gridding, echo classification, and partial-volume effects at plume boundaries. We additionally report that no formal uncertainty propagation was performed and that the metrics are interpreted primarily in terms of relative temporal evolution. The Supplement retains the more detailed algorithmic description and classification examples, but the main text now contains sufficient information to understand and evaluate the Wallangarra analysis independently.

- Several interpretations appear stronger than can be supported directly by reflectivity and polarimetric observations alone. For example, increases in plume volume, reflectivity, or EPC fraction are interpreted as evidence of stronger updrafts, enhanced mass loading, aggregation, larger particles, firebrands, or increased fire intensity. These are physically plausible interpretations, but they are not always uniquely determined by the radar measurements. Thus, I would recommend more moderate casual and physical interpretations in some parts of the manuscript.

R1C4: Thank you for this important comment. We agree that several passages did not sufficiently distinguish between quantities directly observed by radar and their physically plausible, but non-unique, interpretations. We have therefore reviewed the manuscript systematically and moderated causal language throughout. In particular, changes in reflectivity, echo volume, plume depth, and EPC fraction are now described first as changes in radar-observed structure or scattering characteristics. Interpretations in terms of vertical transport, buoyancy, particle loading, particle size, aggregation, firebrands, or fire intensity are now presented as hypotheses or signatures consistent with the observations, rather than as uniquely determined conclusions. We also clarify that reflectivity depends jointly on particle concentration, size, shape, orientation, dielectric properties, phase, attenuation, and sampling geometry. The EPC category has been explicitly redefined as a radar-based class of enhanced and anisotropic pyrometeor echoes; it is not treated as a unique identification of firebrands or as a direct measure of particle mass or fire intensity. Where direct Doppler retrievals or independent observations support a stronger physical interpretation, this is now stated explicitly. We added a new section 6.1 Uncertainties and potential failure modes of radar-derived wildfire diagnostics that describes all limitations from radar observations for all the diagnostics described in each of the previous section and for the demonstration case of Wallangarra.

- The manuscript presents several promising diagnostics, including plume depth, pulsing, rotation, wind-change detection, pyroCu-to-pyroCb escalation, and firebrand-related signatures. However, the uncertainties and potential failure modes are not treated in sufficient depth. This would make the proposed operational pathway more credible.

R1C5: We agree that the uncertainty and potential failure modes of the proposed diagnostics required a more systematic treatment. We have added a dedicated subsection discussing the limitations of plume-depth and pulsing metrics, Doppler and rotation diagnostics, wind-change detection, pyroCu-to-pyroCb escalation indicators, and polarimetric signatures potentially associated with firebrands. The revised text distinguishes direct radar measurements from retrieved or inferred quantities and discusses the effects of radar sensitivity, range, beam geometry, scanning strategy, attenuation, echo-classification uncertainty, velocity aliasing, single-Doppler limitations, availability of clear-air scatterers, and the non-uniqueness of pyrometeor scattering signatures. We also emphasise that no individual radar diagnostic should presently be treated as a deterministic measure or warning criterion. Instead, these products are proposed as confidence-qualified indicators to be combined with satellite observations, lightning data, environmental measurements, fire-behaviour information, and numerical guidance. The associated validation requirements and remaining uncertainties have also been incorporated more explicitly into the TRL assessment and operational roadmap.

- The TRL framework is a useful and original element, but its application is currently somewhat inconsistent (overlapping categories e.g., “5–6” and “6” category). In addition, the manuscript sometimes assigns different readiness levels to hardware, algorithms, deployment procedures, physical understanding, and operational integration. I would recommend either revising or adding a separate TRL assessment level into distinct capability components, for example, radar hardware and mobility, fireground deployment procedures, plume detection, pyrometeor classification, kinematic retrievals, operational integration.

R1C6: We thank the reviewer for this helpful suggestion. We agree that the previous version applied the TRL framework too broadly and combined components with different levels of maturity, which led to overlapping categories and inconsistent interpretation.

We have therefore revised the TRL assessment to distinguish separate capability components: radar hardware and mobility, fireground deployment procedures, plume detection and structural monitoring, low-level wind-change detection, pyrometeor classification and microphysics, kinematic retrievals, multisensor product generation, and operational integration. Each component is now assigned a single indicative TRL based on the level supported by published demonstrations and case studies.

We have also revised the accompanying discussion and Figure 13 to show that these components progress at different rates rather than following a single uniform TRL pathway. The bold outline in each row of the revised figure identifies the current indicative readiness level for that component. This revision removes the previous overlap between categories such as TRL 5–6 and TRL 6 and clarifies the distinction between hardware maturity, scientific understanding, algorithm development, deployment capability, and operational implementation.

- Although I fully understand the concept of "wildfire intelligence", I would recommend the authors to clarify that and provide an explicit definition early in the manuscript.

R1C7: We thank the reviewer for this suggestion. We agree that the term "wildfire intelligence" should be defined explicitly, as it may be interpreted differently across scientific and operational communities. We have therefore added a definition early in the Introduction.

"In this manuscript, wildfire intelligence refers to the timely collection, integration, interpretation, and communication of observational and model information used to support situational awareness and operational decision-making during wildfire incidents. This includes information on fire behaviour, near-fire meteorology, plume development, wind changes, precipitation, and other fire–atmosphere hazards."

Minor comments:

- The Figure 2 caption contains duplicated and apparently incorrect panel labels.

R1C8: Thanks, corrected.

- In Section 4.2, Figure 3a is described in the text as 03:00 UTC, while the caption gives 04:12 UTC. Please correct.

R1C9: Thanks, corrected to 04:12 UTC.

- Figure 4 is hard to read in the current resolution and format.

R1C10: The Figure has been improved: some text has been removed (figure heading), font has been increased and the size of the velocity vectors has been increased.

- Section 4.1 would benefit from a schematic workflow.

R1C11: We thank the reviewer for this suggestion. We agree that Section 4.1 benefits from a clearer visual summary of the processing sequence. We have therefore added a new schematic workflow figure to illustrate how raw polarimetric radar observations are transformed into classified echo fields and subsequently into wildfire-relevant diagnostic products. The figure summarises the main steps, including preprocessing and quality control, feature generation in native polar coordinates, echo classification into pyrometeors, hydrometeors, clear-air returns, and clutter, post-processing, and the derivation of higher-level plume and kinematic diagnostics. We have also revised the text in Section 4.1 to refer explicitly to this workflow and to clarify that echo classification is an intermediate processing layer rather than a final product.

- In Section 4.5, I think that Figure 6 is mistakenly referred to as Figure 4.

R1C12: Thanks, corrected, this is now referred to as Figure 6 in Section 4.5.

- Figure 8 uses different vertical and horizontal scales across panels. These differences should be emphasized to avoid misleading visual comparisons.

R1C13: This is now highlighted in the description of the figure. The former figure 8, now Figure 9 legend now reads:

"Figure 9. (a-f) RHI and horizontal slices at 1 km elevation above ground level of DBZH [dBZ] taken at three chosen timestamps (in UTC) as indicated by vertical dashed lines in figure 8. Red contours indicate pyrometeors and blue contours indicate hydrometeors. Note the different scale on subpanel (e) to the show the full extent of radar observations including the arriving Cumulus Congestus (blue contours) and its interaction with the fire plume (red contours)."

- Could the authors clarify what the blue contours in Figure 8 represent?

R1C14: Thanks, I added:" Red contours indicate pyrometeors and blue contours indicate hydrometeors."

- Figure 9 would benefit from a geographic orientation arrow and clearer explanation of how plume origin was determined.

R1C15: Thank you for this suggestion. We have revised Figure 9 to include a geographic orientation arrow and have expanded both the figure caption and the Wallangarra methods section to explain how the plume-origin reference point was determined. Specifically, the plume principal axis was derived from reflectivity-weighted PCA of the two-dimensional plume footprint at approximately 1 km AGL. Because the sign of a PCA eigenvector is arbitrary, the axis was oriented consistently using the known fire-source and downwind directions and stabilised between successive radar volumes. Plume pixels were then projected onto this oriented axis, and the endpoint corresponding to the minimum along-axis projection was defined as p0p_0p0. This point represents the upwind boundary of the radar-detectable plume footprint at the analysis height, rather than an independent estimate of the exact fire-front or ignition location. The opposite endpoint, p1p_1p1, represents the downwind extent. These clarifications have now been added to Section 5.1 and to the revised Figure 9 caption.

Reviewer 2

General comments

In this paper, the authors address the critical challenges of wildfire intelligence and operational management. They demonstrate the utility of mobile radar monitoring to resolve sub-km plume structure and detect rapid dynamical changes by listing observational requirements and operational constraints, radar-specific requirements and providing a comprehensive review of the main radar observations of smoke plumes and pyroCu/pyroCb convective clouds from prescribed burns and wildfires over the last two decades.

Then, the authors present a new analysis from a recent mobile X-band radar observation of a recent large wildfire in Australia to demonstrate the feasibility to produce portable radar-derived products that could be used by fire agencies. Most interestingly, they develop a road map from research to operations, assess the maturity of the technology for use in operational context to support wildfire management, and make recommendations for future work.

The manuscript is very clearly written and structured and comprehensively addresses challenges in wildfire observations with the stated objective of supporting operational management. Beyond the quasi-exhaustive review of the literature, the work is adequately supported by data, including Supplementary material from the Wallangara wildfire observations, and well suited to NHESS. I therefore recommend publication with minor corrections, as listed below.

Specific comments

135-140 *Tsai et al (2009)* first reported pulsed W-band radar RHI and PPI scans of a prescribed burn smoke plume revealing smoke plume reflectivities in the vicinity of -30 dBZ. *Donnadieu et al (2026)* recently measured a prescribed burn smoke plume in oblique fixed poiting using a miniBASTA FMCW cloud radar (95 GHz). Reflectivities ranged from -40 to -9 dBZ with Doppler +/-7m/s at ranges up to 6 km. Modify fig1 and Table1 accordingly

R2C1: Sentences added, table updated, and references added to the list.

137 Doppler

R2C2: Corrected

148 the condensation level, the lidar attenuated

R2C3: Corrected

161 May et al. (2024) use

R2C4: Corrected

187 fully name AWS

R2C5: Corrected

273 (the so-called Doppler dilemma), making Doppler measurements

R2C6: Corrected

276-281 advantages of RHI and PPI scans are described and combined strategy promoted. However, wildfire observational conditions are often short and changing, pushing operators to chose ; based on their experience, the authors could give their opinion as to the optimum scan strategy in a general case or according to specific observation / smoke behaviour contexts, as done in the following sections on hardware architecture.

R2C7: We added these lines: "Based on the experience gathered during deployments of the UQ-XPOL, we found that the best strategy is often to use sector PPIs with close intervals for the lower tilts, given that the fire plume is often at close range (< 10 km) from the radar. If the radar is positioned along the main axis of the plume, then switching to RHI along the axis would provide exceptional observations as for the Mt Bolton case captured by McCarthy et al. (2019)."

382 "range–height indicator (RHI) at 03:00 UTC": please make timestamps consistent with those displayed in fig.3.

R2C8: Corrected as per R1 similar comment to 04:12 UTC.

Fig.3b what does the vertical line at 04:38 stand for?

R2C9: Thanks, this was an artefact in the radar data and has now been filtered out so it does not appear on the subpanel 3b.

395 the "gradual buildup in low-level reflectivity, followed by a rapid vertical extension, consistent with convective-like burst dynamics": not easily seen from fig.3; any way to show this important feature better (arrow pointing to data, zoomed inset showing different plotting)?

R2C10: Thanks, we added red arrows on the figure 3 and modified the core of the text as well. In the figure legend, it now reads: "Red arrows show low level reflectivity increase followed by rapid vertical extension." In the text, it now reads: "The several-minute spacing may reflect variability in fire heat release, atmospheric stability, plume entrainment, or internal plume dynamics, although the radar observations alone cannot distinguish among these mechanisms. Notably, in some pulses (annotated example in Fig. 4b), an increase in low-level reflectivity appears to precede a more rapid vertical extension of the echo, a pattern consistent with burst-like plume development. Occasional isolated higher-altitude echoes (e.g., around 04:28 UTC) further suggest transient injections of material into elevated layers, possibly associated with localized intensification or interactions with ambient wind shear."

Fig.4 caption: “operational Doppler radar” ; NSW = ? (also line 425)

R2C11: Corrected

415 likewise horizontal wind components and vertical velocity, briefly explain how vertical vorticity shown in fig.4 is obtained

R2C12: We added more information on how vorticity is defined and derived. The text now reads: “Beyond plume depth and reflectivity alone, weather radar can provide physically interpretable diagnostics of the internal dynamical organisation of pyro-convective plumes. Figure 5 illustrates this using the Cooma wildfire case in New South Wales during the 2019–2020 Black Summer wildfires (adapted from Hanstrum et al., 2026), where optical-flow retrievals applied to successive radar images are used to estimate the horizontal motion field and derive associated dynamical quantities, including vertical velocity and vertical vorticity. The arrows in Fig. 5 represent the retrieved horizontal motion vectors, indicating the local direction and magnitude of plume motion inferred from the displacement of radar reflectivity features between successive scans. Vertical vorticity (Fig. 5c) is subsequently calculated from the horizontal velocity gradients as $dv/dx - du/dy$, where (u, m s^{-1}) and (v, m s^{-1}) are the retrieved zonal and meridional velocity components, respectively. Positive and negative values therefore indicate opposite senses of local rotation in the horizontal flow field.

The reflectivity field (Fig. 5a) shows a vertically coherent plume with enhanced echo intensity in the core, while the derived vertical-velocity field (Fig. 5b) highlights concentrated ascent embedded within the broader plume structure. The corresponding vorticity field (Fig. 5c) reveals spatially organised rotational signatures within the convective column. Together with the retrieved motion vectors, these diagnostics illustrate that the plume contained coherent internal circulation rather than behaving simply as a passively advected smoke column.”

Fig.5 caption: “weather radar (Sydney (Terrey Hills) on the 21st of December 2019)” at 06:30 (upper panels) and 07:24 (lower panels).

R2C13: Corrected

509 Fig. 6 documents the evolution of one of the pyroconvective clouds

R2C14: Corrected

572 PCA ?

R2C15: Corrected, Principal Component Analysis named in full in text

Fig.7 caption : indicate meaning of vertical dashed line

R2C16: Corrected. It now reads:

“Vertical dashed lines (at timestamps 024500 UTC, 040400 UTC and 044800 UTC) correspond to the selected timestamps showing radar cross-sections on Figure 9.”

Fig.10 caption: Horizontal slice of…

R2C17: Corrected

Fig.11 This3D representation effort is valuable. To really take adavantage of it, more distinct colour scales for distinguishing better pyrometeors from hydrometeors would be useful, and some length scale for the atmospheric features

R2C18: We thank the reviewer for this useful comment. Figure 11 already uses separate colour scales for the pyrometeor and hydrometeor fields, as indicated by the two corresponding colour bars. We acknowledge, however, that some similar blue/cyan shades occur in both scales and can make the two semi-transparent volumes less readily distinguishable where they overlap. Rather than modifying the original 3-D rendering, we have clarified the distinction between the two fields and their respective colour scales in the figure caption. We have also expanded the caption to provide quantitative information on the spatial dimensions of the displayed domain. The vertical scale is already shown in the figure, while the horizontal extent [and/or grid spacing] is now explicitly stated in the caption, allowing the dimensions of the atmospheric features to be assessed more readily.

691 LES models ?

R2C19: Corrected

691-701 make reference to fig.12

R2C20: Done

Supplementary material

I was unable to find the “full temporal evolution is provided as an animated sequence” in the Supplementary Material: please provide it and make reference to a filename

R2C21: The animation is now include in the supplementary material.

Name/explain the original radar variables DBZH, VRADH, WRADH, ZDR, RHOHV

R2C22: Done

Please provide a reference for the Gaussian Mixture Model (GMM)

R2C23: Done.

Rasmussen, C. (1999). The infinite Gaussian mixture model. *Advances in neural information processing systems*, *12*.

What does the “weighted PCA of the plume mask” stand for?

R2C24: Thank you for highlighting this lack of clarity. Here, “weighted PCA of the plume mask” refers to a principal component analysis applied to the horizontal coordinates of all pixels belonging to the segmented plume footprint. Each plume pixel is weighted by its associated radar reflectivity, so that stronger plume echoes contribute more to the estimated plume orientation while weaker echoes remain included. The first principal component identifies the direction of maximum weighted spatial variance and is interpreted as the plume’s major axis; the second, orthogonal component defines the minor axis. The associated eigenvalues quantify the weighted spatial variance along these two directions. Their ratio is used as a dimensionless measure of plume elongation. We have revised the Supplementary Methods to define this procedure explicitly and to clarify that the equivalent major-to-minor characteristic length ratio is the square root of the eigenvalue ratio.

Fig.S1: subplot labels do not match those in caption: to be harmonized. what do the colours stand for in S1(a)?

R2C25: We added a description for the subplot (a) to explain colour code. We have checked and subplot labels do match those in caption, so we aren’t sure to understand what issue the reviewer is referring to here.

Some contextualization and explanation of fig.S2 is needed in Suppl. Material; also label x-axis

R2C26: Thank you for this suggestion. We have expanded the Supplementary Material to provide additional context and interpretation for Fig. S2. The revised text now explains that the figure shows vertical cross-sections through the Wallangarra plume during its interaction with an approaching precipitating congestus cloud. We describe how the pyrometeor plume and hydrometeor echo differ in their spatial and polarimetric structure and clarify that their partial overlap introduces uncertainty because both particle populations may contribute to the measurements within the same radar resolution volume. We also explain how precipitation-classified voxels were treated in the plume analysis and identify Fig. S2 as an illustration of both the broad-scale separation achieved by the classification and its limitations near the plume–precipitation interface. The Fig. S2 caption has been expanded accordingly, and the x-axis has now been labelled with the appropriate horizontal coordinate and units.