https://doi.org/10.47108/jidhealth.Vol9.Iss5.464
Bediako O, et al., Journal of Ideas in Health (2026); 9(5):1480-1486
Journal of Ideas in Health

Sharper eyes for a rougher sky: a theoretical perspective on how the mathematics of perfectly matched layers support the prevention of turbulence-related injuries to airline passengers and crew

Oscar Bediako[1]*, Debora Alormene[1], Abdul Salam Yussif[2]

Abstract

Background: Turbulence is the leading cause of injury in commercial aviation, and the burden is expected to grow as climate change strengthens the wind shears that generate it. Many such injuries, mostly fractures among cabin crew, follow a potentially preventable sequence in which a hazard is not detected early enough for occupants to be secured. Detection depends on airborne radar and remote-sensing instruments whose design is validated by simulation of electromagnetic wave propagation, and the fidelity of that simulation depends on how faithfully the unbounded atmosphere is represented on a finite grid. This theoretical viewpoint proposes that a specific tool for that truncation, the perfectly matched layer, may be indirectly relevant to passenger and crew safety, and sets out the argument with the underlying mathematics made explicit and accessible.

Methods: We present a self-contained account of the perfectly matched layer as a complex coordinate stretching of the governing wave equations. Using analytic continuation into complex coordinates, we derive its frequency-independent attenuation and the stability, discretization and angular limits that bound its accuracy. The work is theoretical: no human-subject data, clinical outcomes, new simulations or operational aviation-system validation are reported.

Results: The framework shows why the layer is reflectionless in the continuous limit, why its absorption is frequency independent, and precisely where it degrades. These are the properties that govern the accuracy of the sensor simulations on which radar-based turbulence detection partly rests. We also outline a staged pathway for validation against numerical benchmarks, realistic atmospheric data and existing detection methods.

Conclusion: Improved boundary treatment is not a clinical intervention, but may contribute upstream to more reliable hazard detection; any reduction in injuries remains an untested hypothesis. We frame the idea, state its limits honestly, including the radar-invisibility of clear-air turbulence, and invite collaboration between the wave-simulation and aviation-medicine communities.

Keywords: Turbulence, Aviation Medicine, Injury Prevention, Perfectly Matched Layer, Weather Radar, Wave Propagation, Passenger Safety, Clear-Air Turbulence, Cabin Crew, USA

Correspondence: Oscar Bediako (oscar.bediako@unh.edu)
[1]School of Public Health and Allied Sciences, Catholic University of Ghana, Sunyani-Fiapre, Ghana.
How to cite: Bediako O, Alormene D, Yussif AS. Sharper eyes for a rougher sky: a theoretical perspective on how the mathematics of perfectly matched layers support the prevention of turbulence-related injuries to airline passengers and crew. J Ideas Health. 2026 October 31;9(5):1480-1486
doi: 10.47108/jidhealth.Vol9.Iss5.464

Article Info: (**Viewpoint**)
Received: 04 August 2026
Revised: 24 September 2026
Accepted: 25 September 2026
Published: 31 October 2026

© The Author(s). **2026 Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

The Creative Commons Public Domain Dedication waiver (https://creativecommons.org/publicdomain/zero/1.0/) applies to the data made available in this article unless otherwise stated.

Journal Home Page: https://www.jidhealth.com
e ISSN: 2645-9248

Background

Injury in commercial aviation is, to a first approximation, a turbulence problem. National Transportation Safety Board data for United States air carriers show that turbulence accounts for the largest single share of accidents, more than a third of the total across recent decades, and that most of these events produce one or more serious injuries without any damage to the aircraft itself [1]. The injured are disproportionately cabin crew, who are unrestrained while performing their duties, but passengers who are out of their seats or unbelted are also hurt, sometimes severely, with fractures, head injuries and spinal trauma the characteristic pattern [1,2]. The clinical definition of a serious injury in this setting is not trivial: hospitalization beyond forty-eight hours, fracture, internal organ involvement or severe soft-tissue damage [1]. These are real harms to real people, and they are concentrated in a hazard that is, in principle, partly foreseeable. An analysis of National Transportation Safety Board reports for United States commercial flights from 2008 to 2023 identified 136 turbulence accidents with 143 serious and 218 minor injuries; flight attendants were injured in 92.6% of these accidents, most events occurred en-route, and the commonest specified injuries were ankle fractures, leg fractures and spinal injuries [2]. Hospital data tell a similar story: in a series of 43 crew members and passengers treated after eight turbulence

encounters, nearly two-thirds had multiple injuries, extremity and head injuries predominated, about half were admitted and about a quarter required surgery [3]. Turbulence can also be fatal. On 21 May 2024, a Boeing 777 operating Singapore Airlines flight SQ321 over southern Myanmar, probably above an area of developing convective activity, recorded a change in vertical acceleration from +1.35 g to −1.5 g within 0.6 s and an altitude loss of 54 m (178 ft) within 4.6 s; one passenger died and many occupants were injured [4]. Beyond physical trauma, severe encounters can precipitate acute stress reactions and post-traumatic stress disorder, with implications for the fitness of crew to fly and for the wellbeing of passengers [5]. Turbulence injury is therefore at once a trauma-care, occupational-health and public-health problem, and turbulence encounters were the most frequent category of accident in global scheduled commercial aviation in 2023 [6]. For a journal concerned with health, the question of whether such injuries can be anticipated and prevented is a natural one. Two features make this an unusually tractable target for prevention. First, the mechanism of injury is almost always the same, an unrestrained body accelerated against a fixed cabin structure, so that the single protective act of being seated and belted removes most of the risk [1,7]. In the National Transportation Safety Board review of United States air carrier accidents from 2009 to 2018, only one of 123 seriously injured passengers was documented as wearing a seatbelt, and pilot associations continue to advise that passengers keep their seatbelts fastened whenever seated [1,8]. Second, the protective act requires warning: the crew must know, with enough lead time, that rough air is coming, so that the cabin can be secured. The prevention problem is therefore in large part a detection-and-warning problem, and detection is where physics and mathematics enter the clinical story. In the terms of injury prevention, detection and warning act before the injury event, restraint acts during it, and trauma care acts after it; the argument developed here concerns only the first of these. The burden is not static. A consistent body of atmospheric-science evidence indicates that anthropogenic climate change is strengthening the vertical wind shear in the jet streams at cruising altitude, and that the airspace volume containing clear-air turbulence is projected to grow substantially over the coming decades [9,10]. Reanalysis data indicate that this change is already under way: over the North Atlantic, the annual duration of severe-or-greater clear-air turbulence at a typical cruising level increased by an estimated 55% between 1979 and 2020 [11], climate models reproduce reanalysis-based turbulence diagnostics with useful skill [12], and projections suggest that mountain-wave and near-cloud turbulence, not only clear-air turbulence, will become more frequent in a warmer climate [13]. If the hazard is increasing while the cabin countermeasure remains the same belt it has always been, then the value of earlier and more reliable detection rises correspondingly. It is against this background that we make an argument that will at first seem remote from medicine: that an abstract technique from computational wave physics, the perfectly matched layer, may belong among the distal, upstream factors that influence turbulence-injury risk, and that its mathematical properties are worth understanding by anyone interested in why detection sometimes fails. The chain of reasoning is as follows. Convective turbulence and the wind-shear signatures associated with some turbulences are detected, when they are detected, by airborne weather radar and by remote-sensing instruments. Current practice combines several complementary sources: turbulence forecasts derived from numerical weather prediction, increasingly enhanced by machine-learning methods trained on automated eddy dissipation rate reports from aircraft [14,15]; pilot and automated in situ reports; airborne weather radar for convective hazards; and, still at the research and flight-test stage, forward-looking Doppler lidar intended to sense clear-air turbulence ahead of the aircraft [16]. These instruments are designed, sited on the airframe, and their performance predicted, using simulations of how electromagnetic waves propagate, scatter and return [17,18]. Such simulations must represent an effectively infinite atmosphere inside a finite computer, which requires that the artificial edge of the simulated region absorb outgoing waves without reflecting spurious energy back into the scene. A boundary that reflects even weakly corrupts the very quantities, the returned signal, the antenna pattern, the detectable range, that determine whether a hazardous cell is painted on the display in time. The perfectly matched layer is the mathematical construction that makes that boundary, in the ideal case, perfectly absorbing [17]. Its guarantees, and its limits, may therefore propagate, through several intermediate steps that have not yet been quantified, to the reliability of the warning that protects a passenger. We now make those guarantees precise, and afterwards return to what they do, and do not, imply for safety.

Novelty, scope and objectives

The perfectly matched layer is a mature technique, and the derivation below follows established expositions [17,19]. What is new is the explicit interdisciplinary linkage: to our knowledge, the properties of absorbing boundaries have not previously been placed within the causal chain that runs from sensor simulation to turbulence warning and, potentially, to occupant injury, nor have they been presented in a form intended for readers in aviation medicine and injury prevention. The article is a theoretical and conceptual viewpoint. It reports no human-subject data, no clinical outcomes, no new numerical simulations and no validation within an operational aviation system, and any benefit for turbulence detection or injury prevention discussed below is a potential application to be tested, not a demonstrated outcome. Its objectives are threefold: to set out the relevant mathematics accessibly; to state precisely where the guarantees of the method fail; and to propose how the hypothesized link to injury prevention could be validated before any operational or health benefit is claimed.

Methods

Study Design

This is a theoretical, conceptual and mathematical viewpoint. No human participants, human data or animals were involved, and no experimental or clinical results are reported or implied. The mathematics below is a self-contained account of the perfectly matched layer as a complex coordinate stretching; it draws only on results already established in the physics and numerical-analysis literature, which are cited, and follows in particular the complex-coordinate exposition of Johnson [17]. Readers without a background in applied mathematics may wish to read the plain-language overview below and Table 1 first; the technical subsections can be skipped without losing the main argument.

Plain-language overview

A computer can simulate only a finite piece of the sky, so any simulation of a radar signal must stop somewhere. If the edge of the simulated region behaved like a wall, outgoing waves would bounce off it and return as false echoes, much as sound echoes in an empty room. A perfectly matched layer is a thin artificial border, wrapped around the simulated region, that soaks up outgoing waves without creating such echoes, rather like the sound-absorbing foam that lines an acoustic test chamber. The mathematics below explains why such a border can, in principle, absorb waves of every frequency and direction without reflecting them and, equally importantly, the practical conditions under which this ideal is lost. The symbols used are few: u and v are the wave fields, a and b describe the medium, k is the spatial frequency of the wave (its propagation constant), ω its temporal (angular) frequency, σ_x the strength of absorption inside the layer, and θ the angle at which a wave strikes the boundary.

The wave equation and its abstract form

The prototype is the source-free scalar wave equation, $\nabla\cdot(a\nabla u) = b\,\partial^2u/\partial t^2$, for a possibly inhomogeneous medium, with local phase velocity $c = \sqrt{(ab)}$; for lossless propagation a and b are real and positive. It is convenient to split this second-order equation into two coupled first-order equations by introducing an auxiliary field v, so that $\partial u/\partial t = b\nabla\cdot v$ and $\partial v/\partial t = a\nabla u$. Collecting the fields into a state vector w, the system takes the abstract form $\partial w/\partial t = \hat{D}w$. The property that makes this a wave equation is that $\hat{D}$ is anti-Hermitian under an appropriate inner product, which is what produces oscillatory solutions and conservation of energy. Maxwell's equations for electromagnetism, the elastodynamic equations, and the Schrödinger equation can all be written in this same abstract form [17,19], so the construction that follows applies to radar and to any other wave sensor equally.

Radiating solutions

Far from the region of interest, where the medium is homogeneous and unchanging along the truncation direction, radiating solutions decompose into plane waves whose dependence on the truncation coordinate x is the complex exponential exp(ikx). The essential point is that this dependence is an analytic function of x. For energy carried outward the group velocity is positive, and in ordinary media the phase velocity shares its sign, so the propagation constant k may be taken positive.

Analytic continuation into complex coordinates

Because exp(ikx) is analytic, it may be evaluated at complex values of x. If beyond some plane the coordinate is given a positive imaginary part that grows with distance, the exponential separates as $\exp(ik\,\mathrm{Re}\,x)\cdot\exp(-k\,\mathrm{Im}\,x)$, so the field decays exponentially exactly as though it had entered an absorbing material. Crucially, in the interior, where x remains real, the solution is unchanged. The medium is therefore not merely absorbing but reflectionless: interior and exterior are two evaluations of a single analytic function, and where the contour is the real axis, the solution cannot differ from the original. The continued field satisfies the same differential equation, since in the truncation region the equation is invariant along x and the derivative of an analytic function is independent of direction in the complex plane.

The stretch factor and its material meaning

Solving along a complex contour is inconvenient, so we change variables back to the real part. This stretched-coordinate interpretation of the layer is developed in detail by Johnson [17]. Writing the contour as $\tilde{x} = x + if(x)$, every derivative along it becomes the ordinary derivative divided by a stretch factor $s_x = 1 + i\sigma_x/\omega$, where σ_x is an absorption profile and ω the angular frequency. Where σ_x is positive the oscillatory solution decays; where it is zero the equation and its solution are unchanged and there is no reflection. In the scalar equation this stretch factor enters as an effective anisotropic complex material, and for Maxwell's equations it is a theorem that any coordinate transformation is equivalent to the untransformed equations with modified permittivity ε and permeability μ tensors [17]; the uniaxial perfectly matched layer used in most radar solvers is exactly this coordinate stretch absorbed into material tensors [17,18].

Frequency independence

The division of σ_x by ω is deliberate. Under the substitution the wave acquires a decay governed by the ratio k/ω, the reciprocal of the phase velocity; in a dispersionless medium this is fixed, so every wavelength is absorbed at the same rate. Had ω been omitted, short wavelengths would be absorbed faster than long ones and the layer would be spectrally biased, which matters because weather radar and wind-shear sensors operate over finite but non-negligible bandwidths.

Time-domain implementation and stability

In the frequency domain the layer is applied simply by dividing each relevant derivative by its stretch factor. In the time domain the frequency dependence of s_x is handled, in one dimension trivially and in higher dimensions through an auxiliary differential equation that introduces a further field to represent the required integration in time [18]. A well-posed stability analysis of the resulting system, ensuring that the auxiliary equations introduce no growing modes, is a precondition for any use on which safety depends [19].

Limits of the construction

Two limits bound what the layer achieves. First, exactness holds only for the continuous equations; once discretised onto a grid the interface is no longer perfectly reflectionless, and residual reflection appears. This is controlled by turning the absorption on gradually, so that an adiabatic argument drives the reflection down, and by refining the grid [17]. Second, the attenuation scales as $k\cos\theta$, where θ is the angle of incidence, so absorption weakens toward glancing incidence; when radiation originates in a bounded region the incidence angle at a distant boundary is capped and the layer can be made thick enough to absorb the admitted cone [17]. The construction also fails where the medium varies along the truncation direction, since analytic continuation then has no basis [17], although rigorous convergence analysis has recently been extended to some non-uniform configurations, such as wave scattering by step-like surfaces [20]. These are not incidental caveats; they define the conditions under which the mathematics may responsibly be relied upon.

Results

The framework yields five results of a purely mathematical character, none of them experimental or clinical. We state each as a proposition, give the relation on which it rests, and note its bearing on the fidelity of the sensor simulations discussed above. Throughout, σ_x denotes the absorption profile, ω the angular frequency, k the propagation constant along the truncation coordinate, and θ the angle of incidence measured from the boundary normal.

Result 1: Reflectionlessness in the continuous limit

Under the complex stretch $\partial/\partial x \to s_x^{-1}\,\partial/\partial x$ with $s_x = 1 + i\sigma_x(x)/\omega$, a right-going solution $\exp(ikx)$ of the untransformed equation is carried to $\exp(ikx)\cdot\exp[-(k/\omega)\int_0^x \sigma_x(x')\,dx']$. Where $\sigma_x = 0$ the exponent vanishes and the field is pointwise identical to the original; where $\sigma_x > 0$ it decays. The interface reflection coefficient is therefore exactly zero for every frequency, every angle and every polarisation, independent of the profile $\sigma_x(x)$ and even for a discontinuous jump in σ_x. The property is a consequence of analytic continuation alone: interior and layer are two evaluations of one analytic function that agree on the real axis, so no special impedance-matching of a physical material is invoked. For sensing, this establishes that the truncation itself contributes no spurious return; any residual reflection in a computed radar scene must originate in discretisation (Result 4), not in the continuous construction.

Result 2: Frequency-independent attenuation

The round-trip amplitude transmission through a layer of thickness d backed by a perfect reflector is $R = \exp[-2(k/\omega)\int_0^d \sigma_x\,dx']$. Writing $k/\omega = 1/c_x$, the reciprocal of the phase velocity along the truncation coordinate, the attenuation exponent is $-2c_x^{-1}\int_0^d \sigma_x\,dx'$, which contains no explicit factor of ω. In a dispersionless medium c_x is constant, so all spectral components are attenuated equally; had the stretch been defined without the $1/\omega$ factor, the exponent would scale with $k = \omega/c_x$ and short wavelengths would be absorbed preferentially. Because weather radar and wind-shear sensors occupy finite bandwidths, a spectrally flat absorber is a precondition for the simulated return to preserve the shape of a broadband waveform; Result 2 is what guarantees it.

Result 3: Equivalence of the split-field and uniaxial formulations

Absorbing the stretch into the constitutive relations rather than the derivatives replaces the scalar coefficient a by an anisotropic tensor whose diagonal entries carry reciprocal stretch factors, of the form $\mathrm{diag}(s_x^{-1}, s_x, s_x)$ up to the transverse convention, while b acquires the scalar factor s_x. This is the uniaxial (UPML) medium; Berenger's split-field system is the same operator written in characteristic variables. The two are thus not competing schemes but two coordinate readings of a single map, a correspondence that follows for Maxwell's equations from the transformation-optics theorem that any coordinate change is equivalent to modified ε and μ tensors [17,19]. Practically, this shows that the anisotropic complex materials coded into production radar solvers [17,18] inherit, verbatim, the reflectionlessness of Result 1 and the spectral flatness of Result 2.

Result 4: Loss and adiabatic recovery of exactness under discretisation

On a grid of spacing h the continuous dispersion relation is replaced by its discrete counterpart, and the layer interface acquires a non-zero numerical reflection because the discrete and continuous operators no longer share the same analytic solution. Two independent controls drive this residual toward zero. First, grading the profile so that $\sigma_x(x)$ rises smoothly from zero, typically as a quadratic or cubic polynomial across the layer, makes the medium slowly varying on the scale of a wavelength; an adiabatic argument then bounds the reflection, which falls monotonically as the transition is lengthened [17]. Second, refining h brings the discrete operator toward the continuous one, so the reflection of Result 1 is recovered in the limit $h \to 0$. The consequence for sensing is quantitative rather than qualitative: achievable boundary accuracy is set by mesh resolution and layer grading, both of which are design parameters under the modeller's control, so a target reflection floor for a given radar simulation is attainable at known computational cost.

Result 5: Angular degradation at glancing incidence and its bounded regime

For a plane wave incident at angle θ, the relevant propagation constant is the normal component $k_x = k\cos\theta$, so the attenuation exponent of Result 2 scales as $\cos\theta$ and vanishes as $\theta \to 90°$: absorption fails at grazing incidence, and near-glancing components suffer the largest round-trip reflection for a fixed thickness. The regime is nonetheless bounded in the geometry of interest. When all radiation emanates from a compact interior region, the incidence angle at a boundary receding to distance L is capped, and for a cubic computational domain the worst-case angle striking any face approaches $\cos^{-1}(1/\sqrt{3}) \approx 55°$ as L grows, so $\cos\theta$ is bounded below and a finite layer thickness suffices to absorb the entire admitted cone [17]. This delimits precisely the configurations in which a sensor simulation can be trusted and those (near-grazing structure, obliquely truncated features) in which it cannot. Taken together, results 1 to 5 characterize both the guarantees and the failure boundaries of the construction in closed form. Results 1 to 3 establish that, in the continuous setting, the truncation is exact, spectrally flat, and identical across the formulations used in practice; Results 4 and 5 quantify the two mechanisms, discretisation and grazing incidence, by which that exactness is eroded and the conditions under which it is effectively restored. All five are established consequences of the underlying theory; the contribution here is their unified derivation from a single coordinate map and their explicit placement within the chain of reasoning that connects boundary fidelity to the reliability of turbulence warning, and, hypothetically and as yet untested, to occupant injury. Table 1 restates the five results in plain language. No numerical simulation, field measurement or clinical outcome is reported, and none of the five results by itself demonstrates an improvement in turbulence detection.

Table 1. Plain-language summary of the five mathematical results, their relevance to sensor simulation and their main caveats.

Result	Plain-language meaning	Relevance to sensor simulation	Main caveat
1. Reflectionlessness	In the idealized (continuous) mathematics, the absorbing border produces no false echo at all.	The border itself adds no spurious return to a simulated radar scene.	Holds exactly only before the equations are placed on a computer grid.
2. Frequency independence	Waves of all frequencies are absorbed at the same rate.	Broadband radar pulses are not distorted by the border.	Assumes a medium without dispersion.
3. Equivalent formulations	Two common ways of programming the border are the same method written differently.	Results 1 and 2 carry over to the formulations used in practical radar solvers.	Requires a correct and stable implementation.
4. Discretisation	On a computer grid a small false echo appears, but it can be made as small as desired with a finer grid and a gradual border.	Boundary accuracy becomes a controllable design choice with a known computational cost.	Finer grids and thicker borders increase computing time.
5. Glancing incidence	Waves skimming the border at shallow angles are absorbed poorly.	Identifies simulation geometries that can and cannot be trusted.	The border must be thick enough, or far enough away, to cover the expected range of angles.

The results are theoretical properties of the absorbing boundary; they are not measurements of sensor performance or of injury outcome

Discussion

The claim of this viewpoint is deliberately modest and, we hope, defensible. Improved boundary treatment in a wave solver is not a clinical intervention and does not by itself prevent a single injury. What it can do is raise the fidelity of the simulations used to design and evaluate the sensors that provide warning, and warning is the hinge on which turbulence-injury prevention turns. Because the injury mechanism is so consistent, an unrestrained body meeting a fixed structure, and because the countermeasure is so simple, being seated and belted, the marginal value of detecting a hazard a little earlier or a little more reliably is unusually high in this domain [1,7]. Mathematics that makes sensor simulation more trustworthy could therefore, at one remove, contribute to passenger and crew safety, although the size of any such contribution is unknown. Several qualifications are essential, and stating them plainly is part of the argument rather than a retreat from it. The most important is that a large and growing share of the turbulence that injures people is clear-air turbulence, which produces no significant radar return and cannot be detected by conventional weather radar at all [9]; for this hazard the relevant instruments are forecasting models and forward-looking optical techniques such as LIDAR [14,16], and no improvement in electromagnetic boundary treatment will make invisible air visible. In the National Transportation Safety Board analysis, convective turbulence accounted for 57% and clear-air turbulence for 28% of turbulence injuries [1,8], so the share of the hazard that radar can in principle address is substantial but far from complete. The honest scope of the present argument is therefore the radar-detectable portion of the hazard, principally convective turbulence and certain wind-shear signatures, together with any optical or radio sensor whose design is itself validated by wave simulation. Within that scope, the recurring finding in accident investigations, that a convective cell was not painted as a significant return until too late, is precisely the kind of failure that higher-fidelity sensor and propagation modelling might help to reduce, although no such effect has yet been demonstrated. A second qualification is that the perfectly matched layer inherits the limits derived above: it is exact only in the continuous limit, weakens at glancing incidence, requires a stable time-domain implementation, and fails in media that vary along the truncation direction [17,19]. Any safety-relevant claim must be verified against these limits and against recognized aviation certification standards, not assumed from the mathematics. A third is that detection is only one link in a longer chain that includes forecasting, dissemination, crew procedure and the passenger's own choice to remain belted; the National Transportation Safety Board has repeatedly emphasized these operational and human factors [1], and better sensing complements rather than replaces them. With those qualifications in place, the constructive proposal is simply that the wave-simulation community and the aviation-medicine community have more to say to one another than either currently supposes. The mathematics of reflectionless truncation is mature and general; the epidemiology of turbulence injury is well characterized and, unusually, points to a single dominant and partly detectable hazard with a single simple countermeasure [1,2]. Bringing the two together, so that improvements in simulation fidelity are evaluated explicitly in terms of their contribution to timely warning and thence to injury reduction, seems to us a worthwhile and largely unexplored idea. The staged approach outlined below offers one way of testing it.

Pathway to validation

Because the link between boundary treatment and injury prevention is hypothesized rather than demonstrated, we propose that it be tested in four sequential stages, each of which must succeed before the next becomes meaningful and before any operational or injury-prevention benefit is claimed.

Stage 1, numerical benchmarks: An implementation should first reproduce problems with known answers: the analytical reflection coefficient of Results 1 and 2 in one dimension; the convergence of residual reflection with grid spacing and profile grading predicted by Result 4, using established adiabatic-absorber benchmarks [17]; long-time stability tests for anisotropic and dispersive media [19]; and grazing-incidence test cases [17]. Results should be compared with independently developed solvers and with large-domain reference solutions in which the boundary is placed so far away that its reflections cannot reach the region of interest within the simulated time [18,19].

Stage 2, realistic atmospheric scenes: Simulated radar returns should then be computed for realistic three-dimensional fields of reflectivity, wind and turbulence intensity taken from high-resolution numerical weather prediction or reanalysis, including documented convective encounters such as that of flight SQ321 [4], and compared quantitatively with simulations that use conventional absorbing boundaries. Appropriate outcome measures are the error in the simulated returned power and Doppler spectrum width, the modelled detection range for hazardous cells and the computational cost.

Stage 3, comparison with existing detection methods: Any improvement in simulated sensor performance should be assessed against what operational systems already achieve, using automated in situ eddy dissipation rate reports from aircraft as ground truth [14] and comparing the probability of detection, false-alarm ratio and warning lead time with those of current airborne radar turbulence algorithms, forecast products, including machine-learning forecasts [14,15], and experimental lidar systems [16].

Stage 4, operational and injury outcomes: Only if the earlier stages show a meaningful gain in warning lead time or reliability would it be appropriate to examine operational and health outcomes, for example the interval between hazard detection and illumination of the seatbelt sign, and turbulence injury rates per departure derived from National Transportation Safety Board and airline safety data [1,2]. Such studies would need to account for confounding by crew procedures, route and season, and for the long-term rise in turbulence frequency [11].

Limitations of study

This work has important limitations. First, it is theoretical and conceptual: it contains no human-subject data, no clinical or injury outcome data, no new numerical simulations and no validation within an operational aviation system, so its conclusions are hypotheses rather than findings. Second, the causal chain from boundary fidelity to sensor design, to warning, to cabin preparation and finally to injury has several intermediate links, none of which has been quantified for this purpose; in practice, the contribution of absorbing-boundary errors to overall detection performance may be small compared with other sources of error, such as sensor physics, signal processing, atmospheric variability and human factors. Third, clear-air turbulence, which accounts for a substantial share of injuries, lies beyond the reach of radar and is addressed here only insofar as optical or other sensors are themselves designed with the help of wave simulation. Fourth, the mathematical results are established ones, and their applicability depends on the assumptions stated in the Methods, notably a medium that is uniform along the truncation direction, adequate grid resolution and a stable implementation. Finally, studies linking turbulence detection to injury outcomes remain scarce, which limits the extent to which any upstream improvement can currently be translated into an estimate of injuries prevented.

Conclusion

Turbulence is the principal cause of injury in commercial aviation, its burden is projected to grow, and the injuries it causes are, for the detectable portion of the hazard, potentially preventable given timely warning. Warning depends on sensors whose design rests on wave simulation, and the fidelity of that simulation rests in part on how the unbounded atmosphere is truncated onto a finite grid. We have set out, with the mathematics made explicit, why the perfectly matched layer provides a principled and in the ideal case perfect solution to that truncation, and precisely where it does not. We do not claim a measured reduction in injury, and we are explicit that clear-air turbulence lies outside the reach of radar and hence of this argument. What we do claim is that the mathematical soundness of the absorbing boundary is a plausible, if distant and as yet unquantified, upstream influence on how reliably a hazard is seen, and therefore potentially on the health and safety of those on board. This proposition requires the staged validation described above before any benefit for turbulence detection or injury prevention can be claimed. We offer the idea in the hope of prompting collaboration across the disciplines it spans.

Abbreviation

CAT: Clear-Air Turbulence; LIDAR: Light Detection and Ranging; NTSB: National Transportation Safety Board; PML: Perfectly Matched Layer; UPML: Uniaxial Perfectly Matched Layer.

Declaration

Acknowledgment

The authors thank the anonymous reviewers for their constructive comments, which improved the manuscript.

Funding

None.

Availability of data and materials

Data sharing is not applicable to this article, as no datasets were generated or analyzed; the work is theoretical. Correspondence may be directed to the corresponding author: oscar.bediako@unh.edu.

Authors' contributions

The corresponding author conceived and directed the study, developed the mathematical framework, and drafted and revised the manuscript. All authors have read and approved the final article.

Ethics approval and consent to participate

We conducted the research in accordance with the Declaration of Helsinki. Ethics approval is not applicable. This viewpoint did not involve human participants, human data, or animals. Ethical approval and informed consent were therefore not required.

Consent for publication

Not applicable

Competing interest

The authors declare that they have no competing interests.

Open Access

This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The Creative Commons Public Domain Dedication waiver

(http://creativecommons.org/publicdomain/zero/1.0/) applies to the data made available in this article unless otherwise stated.

Author Details

[1]Department of Mathematics and Statistics, University of New Hampshire, USA.

[2]Department of Mathematics and Statistics, Youngstown State University, USA

References

1. National Transportation Safety Board. Preventing turbulence-related injuries in air carrier operations conducted under Title 14 Code of Federal Regulations Part 121. Safety study SS-21/01. Washington (DC): NTSB; 2021.
2. Shekhar AC, Ruskin KJ. Injuries due to in-flight turbulence during United States commercial airline flights (2008-2023). Am Surg. 2026;92(5):1398-403. https://doi.org/10.1177/00031348251397042
3. Suriyawongse T. Air turbulence associated with crew member and passenger injuries on commercial aviation. Bangkok Med J. 2021;17(1):20-3. https://doi.org/10.31524/bkkmedj.2021.11.004
4. Transport Safety Investigation Bureau. Transport Safety Investigation Bureau preliminary investigation findings of incident involving SQ321. 2024 May 29. In: Ministry of Transport Singapore [Internet]. Singapore: Ministry of Transport. Available from: https://www.mot.gov.sg/news/Details/transport-safety-investigation-bureau-preliminary-investigation-findings-of-incident-involving-sq321
5. Vuorio A, Bor R, Budowle B, Gray A, Suhonen-Malm AS. Assessment policy of post-traumatic stress disorder in aviation and its practical application using turbulence-triggered trauma as an example. Front Public Health. 2025;13:1505004. https://doi.org/10.3389/fpubh.2025.1505004
6. International Civil Aviation Organization. Aviation safety improved despite the challenges of pandemic recovery. 2024 Aug 29. In: ICAO Newsroom [Internet]. Montréal: ICAO. Available from: https://icao.int/Newsroom/Pages/Aviation-safety-improved-despite-the-challenges-of-pandemic-recovery.aspx
7. Federal Aviation Administration. Preventing injuries caused by turbulence. Washington (DC): FAA; 2022. Available from: https://www.faa.gov/newsroom/preventing-injuries-caused-turbulence
8. International Federation of Air Line Pilots' Associations. Turbulence injury mitigation. Briefing leaflet 25AAPBL03. 2025 Nov 24. In: IFALPA [Internet]. Montréal: IFALPA. Available from: https://ifalpa.org/media/4341/25aapbl03-turbulence-injury-mitigation.pdf
9. Williams PD. Increased light, moderate, and severe clear-air turbulence in response to climate change. Adv Atmos Sci. 2017;34(5):576-86. https://doi.org/10.1007/s00376-017-6268-2
10. Storer LN, Williams PD, Joshi MM. Global response of clear-air turbulence to climate change. Geophys Res Lett. 2017;44(19):9976-84. https://doi.org/10.1002/2017GL074618
11. Prosser MC, Williams PD, Marlton GJ, Harrison RG. Evidence for large increases in clear-air turbulence over the past four decades. Geophys Res Lett. 2023;50:e2023GL103814. https://doi.org/10.1029/2023GL103814
12. Williams PD, Storer LN. Can a climate model successfully diagnose clear-air turbulence and its response to climate change? Q J R Meteorol Soc. 2022;148(744):1424-38. https://doi.org/10.1002/qj.4270
13. Kim SH, Kim JH, Chun HY, Sharman RD. Global response of upper-level aviation turbulence from various sources to climate change. npj Clim Atmos Sci. 2023;6:92. https://doi.org/10.1038/s41612-023-00421-3
14. Muñoz-Esparza D, Sharman RD, Deierling W. Aviation turbulence forecasting at upper levels with machine learning techniques based on regression trees. J Appl Meteorol Climatol. 2020;59(11):1883-99. https://doi.org/10.1175/JAMC-D-20-0116.1
15. Zhu C, Ng KKH, Ban W. Enhancing clear-air turbulence forecasts: a WTCNN-LSTM framework utilising energy dissipation rate. In: AIAA AVIATION Forum and ASCEND 2025; 2025 Jul 21-25; Las Vegas (NV). Reston (VA): American Institute of Aeronautics and Astronautics; 2025. https://doi.org/10.2514/6.2025-3715
16. Vrancken P, Herbst J. Aeronautics application of direct-detection Doppler wind lidar: an adapted design based on a fringe-imaging Michelson interferometer as spectral analyzer. Remote Sens. 2022;14(14):3356. https://doi.org/10.3390/rs14143356
17. Johnson SG. Notes on perfectly matched layers (PMLs); 2021. Preprint. Available from: https://arxiv.org/abs/2108.05348
18. Liu K, Huang T, Zheng L, Jin X, Lin G, Huang L, et al. Efficient parallel FDTD method based on non-uniform conformal mesh. Appl Sci. 2024;14(11):4364. https://doi.org/10.3390/app14114364
19. Duru K, Kreiss G. The perfectly matched layer (PML) for hyperbolic wave propagation problems: a review; 2022. Preprint. Available from: https://arxiv.org/abs/2201.03733
20. Lu W, Zheng W, Zhu X. Perfectly matched layer method for the wave scattering problem by a step-like surface. SIAM J Numer Anal. 2025;63(2):744-71. https://doi.org/10.1137/24M1654221.