10 May 2026

Response to RC1 & RC 2 – GMD 2026-429

***Response to RC1***

*We thank the reviewer very much for his positive, detailed and valuable contribution to improve the manuscript. We have addressed the feedback in the revised version of the manuscript and respond to all comments below.*

**General comments**

The authors present a new interesting tool, AgPaDS, developed to rapidly simulate Lagrangian transport in the atmosphere, coupled with advanced visualization and interactive simulation configurations, for agricultural applications. They compare its performance with HYSPLIT, one of the mostly used atmospheric transport models, demonstrating a significant increase in computational efficiency across different tests, thanks to the use of GPU. This tool goes into the direction of solving persistent challenges faced by agricultural researchers, such as (a) the simulation time (which is dealt by directly in the model) and (b) the coupling of windborne pathogen transport with plant epidemiological dynamic models (which the authors propose as a possible advancement for this tool).

**Specific comments**

Undoubtedly, the improvement in computation time compared to the benchmark is outstanding, and the visualization tools are also remarkable.

- *Thank you for the positive feedback.*

The article's focus is more on the algorithmic aspects than on the Lagrangian implementations, for which the authors drew inspiration from NAME. I recognize that this is a very technical paper, and as someone not deeply familiar with Lagrangian modeling at this level of implementation, I admit it is not always easy to follow.

For instance, (i) I would have preferred the model description (section 3) to appear first in the Materials and Methods;

- *Before submission, we had also discussed the ordering of sections 2 and 3, as both, the current ordering and a swap, have their advantages/disadvantages. The feedback is useful for us, as it indicates that, indeed, a swap in order would make it easier for readers to follow the manuscript. We have addressed the comment by inter-changing sections 2 and 3 in the revised manuscript.*

- (ii) acronyms (GPU, ECMWF, CUDA) are used without explanation.
    - *We have inserted the definition of all acronyms on first mention.*

On the other hand, despite the claims of a better compatibility of this tool with crop or epidemiological models, the coupling does not seem straightforward.

- *Our main argument is not a better compatibility in the sense of an easier or more user-friendly loose coupling of AgPaDS with epidemiological models (i.e. use file output of AgPaDS as input for epidemiological model), but instead a new possibility for tight coupling of the two modelling approaches, i.e. execution of atmospheric transport simulations as part of the main time-loop of a spatiotemporal epidemiological model, enabled by performance gains and model architecture (e.g., gridded release). We have already implemented this in a meteorology-driven SIR meta-population model (unpublished) and have conducted initial tests showing promising results. Originally, we had planned to include this in the current manuscript, but as it represents another layer of model development and testing and the manuscript is already comprehensive, it is beyond the scope of this paper. However, to address the* *comment, we have included an additional paragraph at the end of the discussion to summarize the approach in conceptual terms. In a follow-up paper, we will introduce details.*

I commend the authors for their accurate model evaluation setups, both in terms of experimental design and in terms of indicators used to measure the comparisons.

- *Thank you very much.*

However, I believe that a couple of points require revision or at least further discussion.

In the second experiment (Table 3), the authors state that differences between HYSPLIT simulations and AgPaDS are less than one order of magnitude, but the data actually show a one-order-of-magnitude difference; this is not inadequate per se, but needs to be better contextualized.

- *Corrected. Considering all metrics it is less than or equal to approximately one order of magnitude. We agree with the second part of the comment, indeed it is important to consider typical values and uncertainties regarding this quantity. This is included in the discussion.*

Moreover, in the third set of experiments (~Line 735), when comparing the simulation of atmospheric transport of *Phakopsora spores* by hurricane Ivan if it would be feasible to compare the simulated and observed deposition/presence of the soybean rust infections in USA (and not on the severity), instead of leaving it qualitatively.

- *We are not sure if we understand the comment correctly. It sounds as if the reviewer is suggesting that we are comparing the severity of soybean rust infections in qualitative terms with simulated deposition patterns, which we are not. As illustrated in Figure 7, we are comparing simulated deposition patterns with locations of sites of first detection of the pathogen in the USA. Unfortunately, a more exact quantitative evaluation is not possible given the available data, because the source strength is unknown (no direct measurements of pathogen emission rates nor sufficient information about infected area and infection intensity to estimate source strength in Colombia). Please note that, here we directly follow the authors approach of the original study – Isard et al – who have stated the same, the empirical data is too sparse to quantify emission and deposition rates in more detail, so they consider unit release in their simulations.*

My last question is, given the strong inspiration drawn from NAME, is there a specific reason why the authors chose to benchmark against HYSPLIT rather than on NAME?

- *As part of the development and testing phase, we have drawn inspiration not only from NAME, but also from IAMS, HYSPLIT, MPTRAC, and others (as referenced in the text). The decision to benchmark against HYSPLIT was motivated mainly by: (i) HYSPLIT is available online for everyone without license requirements; (ii) HYSPLIT is very widely used, including in directly related studies in crop epidemiology (e.g., Radici et al, collaboration of NOAA with FAO on desert locusts).*

**Minor corrections**

- L 15: GPU, which is the Graphical Processing Unit.
  - *Included explanations of acronyms throughout.*
- 38-39 “et al.” is sometimes italicized, but not always. Please homogenize throughout the text.
  - *Corrected, thanks.*
- 40. So far, the introduction quotes very general papers, without going into details. For instance, the authors do not tell the name of any of these “devastating crop diseases” (L. 38) and seem to suggest that

long-distance-dispersed plant pathogens cause 17-30% of yield loss (L.40), while these percentages aggregate all infectious plant diseases.

  - *The introduction is kept fairly general on purpose, as the manuscript is not about one or two specific diseases, but about a new model that can be applied to a range of airborne diseases, and bioflows more generally. However, we agree that it could profit from some more specifics, and have included concrete examples. It was not our intention to make it seem as if airborne diseases alone cause these losses. We have rephrased to avoid misunderstandings.*

- 52: atmospheric transport models or Atmospheric Transport Models (as in L. 10)?
  - *Have ensured consistency throughout.*
- 80: What does I/O mean? Please, define acronyms throughout the text.
  - *Input/Output; have included explanation of all acronyms throughout.*
- 85: “These challenges evolve around, unknowns around pathogen viability decay during atmospheric transport: uncertainty estimates for processes involved in atmospheric transmission of crop pathogens”. What does it mean?
  - *Rephrased to clarify. One key example of unknowns is a lack of exact experimental data on the proportion of spores, out of large populations of fungal spores (millions/billions), that survive certain regimes of meteorological factors like UV, temperature and humidity. One key example of previous uncertainty estimates in the literature is the work of Aylor, which is referenced in the sentence.*
- L81-132: these bullet points are interesting and cover exhaustively the issues with ATMs, but it think they could gain in readability – for example, by splitting or shortening some sentences and capitalizing the initial letter of each point.
  - *Noted and revised in attempt to improve legibility.*
- 97: I wonder if the authors meant “links” and not “vertices” (which I assume is asynonym of “nodes”).
  - *Yes, corrected, it is links / edges.*
- L 120: Missing full stop.
  - *Corrected, thanks.*
- 140: Is the figure correctly placed here?
  - *The figure is placed on top of the page on which it is first referenced.*
- 210: P in Python should be capitalized, here and in the rest of the text
  - *Done; thanks.*
- 219: Specify what ECMWF is.
  - *Acronyms defined on first mention.*
- 361-369: What do subscripts v, p, lambda and phi stand for? Also, is m_v the same as m_i (I do not think so)?
  - *The subscripts lambda, phi, p denote spatial dimensions. In the initial version of the manuscript these are introduced in a later section of the manuscript, but they should have been introduced upon first mention, so this is corrected now. The two different notations m_v and m_i are chosen to represent different quantities. As explained in the text, m_v is used to denote a function to*

*model viability decay, in the equation introducing the conceptual model, whereas m_i is used to denote the material carried by each simulation particle, i.*

- 394: SSD?
  - *Acronyms are now defined at first mention.*
- 397: CUDA? Quoted 33 times, also in Fig. 2
  - *Acronyms are now defined at first mention.*
- 460-461: I am not sure of the rationale. Is N the number of simulation particles as in L. 371?
  - *Yes, once introduced, the same notation is used throughout.*
- 490: Consider telling more explicitly the reader the dimensions of wind data in ERA 5 (x, y, pression), not just as unit of measures.
  - *This has been explained in the text already (see section 3.4.1 on model grid and section 3.4.3 on input data.*
- 503: Consider telling the reader what a timescale is.
  - *Included short summary of Lagrangian timescale.*
- 510: no space before comma.
  - *Checked for spaces before comma.*
- 537-554: A very interesting overview of possible techniques of implementing viability decay!
  - *Thanks.*
- 570: This is not the first time I see this value of this parameter used for wet deposition. This is quite evidently computed on an approximation, since 25.4 mm = 1 inch and 63.2% is just 100*(1 – exp(-1)), or the expected loss after 1 unit of time. It is a pity that there are better estimations for this parameter.
  - *As noted in the text, we choose to follow the parsimonious and simple, yet useful, parameterization for wet deposition as implemented in IAMS. There are various other parameterizations of wet deposition in the literature on atmospheric dispersion modelling, but few for the specific case of crop pathogens. We choose here a proven generic approximation, noting that it is straightforward to adapt this part of the model to specific crop pathogens in future use-cases, e.g. when specific empirical data are available for specific pathogen types.*
- 675: Despite atmospheric trajectories and material deposition look very similar, the same can not be said about Lagrangian particles in the air. Why?
  - *Whilst on first sight it appears like there may be much larger differences between the Lagrangian particles in air compared with the deposition patterns and trajectories, this is mostly a consequence of the type of visualization in the figure - the 2D point cloud shows 10 K simulation particles, which, in the central parts of the plume, overlap strongly, whilst in the outer parts individual particles / smaller subsets of particles follow slightly different paths. We choose to visualize individual Lagrangian particles instead of particle densities, as we want to illustrate a core numerical characteristic of the model, discrete simulation particles. The fact that small subset of Lagrangian particles in air deviate is to be expected, as these are stochastic particle ensembles. The gridded deposition plots show that the overall distribution of particles is indeed in good agreement. Trajectories are deterministic, so differences here are not due to random variations, but other model differences, e.g. interpolation scheme of meteorological data to trajectory position.*

- 692: Also relative differences HYSPLIT – AgPaDS look quite important int Table 2. Could you please discuss?
    - *There are differences between the two models, but this is to be expected, and is discussed in detail in the manuscript (see second paragraph, section Discussion).*
- 710: The same can be said of Table 3, for example of the median deposition value. Despite what you state in L. 712, thz difference between $5.8* 10^{-10}$ and $2.8* 10^{-11}$ is ~1 order of magnitude.
    - *Corrected to smaller/equal to approximately one order of magnitude. Importantly, as stated above by the reviewer and already noted in the discussion, the extent is acceptable considering given uncertainty ranges, data sparsity, and the main modelling goal here of obtaining a highly performant yet sufficiently good approximation.*
- 735: "Whilst the available data does not allow for exact quantitative evaluation". This sentence makes me ask if (1) there were no better episodes of LLDto test the model performances, such as stem rust of wheat in East Africa or Asia or (2) a test to compare the presence (and not on the severity) of the soybean rust infections in the USA and your simulations would be feasible.
    - *The phrasing "exact quantitative evaluation" is intended to refer to an experimental design as used in large-scale tracer experiments, including the exact measurement of emission rates at the source along with measurements of spore concentration in air as well as deposition to the surface at various locations and times around the source. To the knowledge of the authors, on landscape to global scales this type of data does not exist for crop pathogens. We have added a clarification on the phrasing. The core objective is the development of a broadly applicable simulation tool, rather than a specific model for a specific disease outbreak like stem rust in East Africa. Therefore, we chose to evaluate against 2 existing well-tested simulation frameworks, complemented by a case-analyses that contains what the reviewer is suggesting, a comparison to presence/absence by looking at first detection sites in the USA. In future studies, we plan to conduct further evaluations as part of specific use-cases.*
- Line 873: "in comprehensive".
    - *Corrected, thanks.*

**Citation**: https://doi.org/10.5194/egusphere-2026-429-RC1

10 May 2026

Response to RC1 & RC 2 – GMD 2026-429

***Response to RC2***

*We greatly appreciate the positive, constructive and detailed feedback of the reviewer, which has helped improving the manuscript. We have addressed the feedback in the revised version of the manuscript and respond to all comments below.*

This manuscript presents AgPaDS, a GPU-accelerated Lagrangian atmospheric transport model with interactive 3-D in situ visualisation, motivated primarily by applications in crop epidemiology. The paper describes in exceptional technical detail the software architecture, GPU implementation, and visualisation capabilities, and demonstrates that the model reproduces results broadly consistent with established atmospheric transport models (HYSPLIT, IAMS) while achieving substantial performance gains. This is a strong and valuable model description paper but I have some concerns regarding balance, framing, and evaluation that should be addressed prior to publication.

- *Thank you very much for the positive feedback, evaluating the manuscript as a strong and valuable model description paper with exceptional technical detail. We are happy to hear that the core focus of the paper is perceived well. We address the remaining concerns and all comments below.*

Major issues

1. The work represents a significant engineering and visualisation achievement and fulfills many criteria of a GMD model description paper. However, while the manuscript is positioned as a model description paper, a substantial fraction of the text focuses on GPU architecture, CUDA/OpenGL implementation, GUI design, and visualisation pipelines, often at greater length and detail than the atmospheric or biological modelling itself. Hence, the manuscript would benefit from clearer framing of its scientific contribution, greater balance between software engineering and atmospheric/biological modelling, and a more critical discussion of limitations and use-case constraints because the scientific novelty risks being overshadowed by the engineering narrative. Consider shortening the discussion of enabling technology and reallocating space to sensitivity of results to modelling assumptions, limitations and failure modes, and scientific implications for crop epidemiology beyond "speed + visualisation".

- *Thank you very much for acknowledging the work as a significant engineering and visualization achievement that fulfills many criteria of a GMD model description paper.*
- *We understand there are different perspectives regarding the balance between technical/methodological focus on the one hand and atmospheric/biological modelling on the other hand. Considering that - at the core - the paper is about technical/methodological innovation, we remain convinced that it is important and appropriate for this manuscript to include technical model description details. The technical focus of the paper has been clarified in the abstract and in the last paragraph of the introduction section. In addition, we have re-structured the manuscript, swapping sections 2 and 3, which places more importance on atmospheric and biological models by placing these more prominently in the overall article. Also, we have extended parts of the discussion.*
- *We find the framing of the scientific contribution clearly formulated in the original manuscript, see at the end of the introduction section, and also at the beginning of the model goals section and sections on test cases and discussion. We would therefore not like to change the text.*

2. Model evaluation is largely based on qualitative visual agreement, relative comparison to HYSPLIT and IAMS, with broad consistency against empirical soybean rust observations. While this is reasonable given limited data, agreement with another model does not constitute validation, and some reported differences (e.g. 24% longer mean trajectories; 34-43% differences in plume spread) are non-negligible for epidemiological applications. The evaluation results should be clearly framed as benchmarking against existing models, not validation. Also consider strengthening the discussion of which deviations matter (and for which questions), which applications are robust to observed discrepancies, which

applications require caution or additional calibration. I also recommend adding idealised test cases (e.g. homogeneous flow, no turbulence) where expected behaviour is conceptually clear and demonstrated.

- *Model testing is clearly framed as evaluation against existing models, and not as validation in strict sense (see original and revised manuscript). The term validation is also not used.*
- *The initial manuscript already includes a discussion of advantages and disadvantages of the modelling framework and use cases, which we have extended in the revised manuscript, along with a dedicated table (S4) discussing feasibility of different use-cases that include bottlenecks.*
- *Overall, the manuscript strikes a balance between technical model description details, domain-specific modelling approaches, and model testing that includes a comprehensive comparison with two widely used ATM modelling frameworks. Additional idealized test cases and extensive discussion of implications of deviations for different types of use-cases and questions are beyond the scope of this manuscript. However, in ongoing work, we are focusing on more applied use-cases, extending the testing and including disease surveillance data.*

3. Epidemiological relevance is aspirational rather than demonstrated because the manuscript does not present a clear end-to-end epidemiological analysis or a decision-relevant use case. This means that most epidemiological benefits are presented as future potential, rather than demonstrated outcomes of the new modelling framework. No crop-disease risk metric is produced, no comparison to observed disease timing/intensity beyond spatial coincidence is provided, and no assessment is made of how interactive visualisation could change decisions. I recommend that the authors either add one concrete demonstration of epidemiological utility or more explicitly limit and qualify claims regarding operational relevance and early-warning systems.

- *We do not think that epidemiological relevance is merely aspirational. The focus of the manuscript is on method development and the design of a new method that enables addressing key open challenges in contemporary epidemiology research, as described in this manuscript, has epidemiological relevance. Further, we have included one test-case (soybean rust) that shows epidemiological utility, and have chosen the test locations for the comparison with HYSPLIT such that these provide realistic outbreak locations on wheat producing regions considering known previous infection sites. Also, we contextualize our work based on previous literature, show by model formulation and testing, and discuss epidemiological relevance in detail (see e.g. table S4). However, we understand there are different legitimate perspectives on terminology of epidemiological relevance, and fully acknowledge that the manuscript does not describe an epidemiological model application study, as e.g. the work of the reviewer assessing the effect of green bridges for long range dispersal of wheat rusts using an existing ATM.*
- *During initial conceptualization of the manuscript, we had considered adding an additional epidemiological case-study but it became evident that this is beyond the scope of the current manuscript, which is already comprehensive, covering the core focus on method development and testing in detail. In the revised manuscript, we have added a paragraph at the end of the discussion to clarify the linking of the atmospheric transport model introduced here with epidemiological models and to provide a glimpse at ongoing and future work.*
- *We are not sure we understand which claims regarding operational relevance and early warning systems the reviewer seems to perceive as overstated. It is not our intention to overstate operational relevance, and we have therefore already in the original manuscript chosen careful phrasings (e.g., as a "model aim", as part of "discussion of feasibility" and as part of "potential for future work").*

4. The atmospheric transport formulation closely follows established Lagrangian frameworks, and many parameterisation schemes are largely adopted from earlier work. The rationale and sensitivity for those parameter values, and the interaction with GPU performance constraints are not fully discussed. A short sensitivity analysis or explicit discussion of which parameters are intended to be exploratory versus physically constrained would substantially strengthen the model description.

- *We agree with the reviewer regarding clarity of simulation parameters and have strengthened the model description and discussion by adding a new table to the SI (table S1) with selected key parameters, including common ways for process- and data-based parameter definitions and options for interactive parameter configurations.*
- *Regarding the other comments, we are not sure we understand. In our view: (i) the rationale for model formulation and design is provided and discussed in detail (e.g., introduction with specific literature context that provides rationale; model goals; model description; discussion of advantages/disadvantages and feasibility of use-cases); (ii) sensitivity to some selected parameters/sub-models is included (e.g., release height in test cases, viability decay model variants), and a full sensitivity sweep for all parameter values is clearly beyond scope; (iii) the effect of different sub-model choices and parameter values on computational performance is already discussed (see e.g. section on computational performance and summary of computational implementation in model description)*

5. While I am not a specialist in high-performance computing, the validity of the computational speed comparison to HYSPLIT is not entirely clear, as performance is evaluated against a single-CPU configuration rather than commonly used parallel CPU setups. Whilst this fact is acknowledged, it should also be stated that such configurations would reduce the apparent advantage of AgPaDS. Clarifying this limitation explicitly would strengthen the transparency of the performance assessment rather than weaken the contribution. In addition, while the manuscript draws conceptually on NAME (e.g. for turbulence parameterisation), no explicit comparison to NAME is presented. This is not necessarily required, but the absence should be more clearly acknowledged and its implications discussed, particularly in the context of performance claims and plume spread differences.

- *The publicly available HYSPLIT version that is used here as baseline for computational performance estimates is one of the most widely used ATMs. We have added a brief discussion to motivate its use for model testing.*
- *In the discussion of the original manuscript, we had already included an acknowledgement of expected performance differences regarding HYSPLIT, along with contextualization to other existing modelling approaches, including a recent GPU implementation.*
- *Further to this, we have added to the discussion in the revised manuscript that existing multi-CPU setups provide speed-ups, noting that this does not change the main argument, as (a) the performance gains achieved here by the massively parallelized GPU implementation can be expected to be higher than those in multi-CPU setups on individual workstations for very large particle numbers (because of the limited number of CPUs compared with processing units on GPU, the efficient implementation shown here, e.g. scaling of compute time with number of Lagrangian particles), (b) importantly, the interactive in-situ visualization methods introduced here are innovative and not currently feasible with existing ATMs; (c) as noted in the discussion, a certain performance gain is not the main goal, but rather an efficient implementation with novel features and flexibility for future domain-specific customizations.*

Minor issues

1. The manuscript is very long and there are several sections that repeat the same points, particularly in the Introduction and the Discussion. Careful revision of long sentences and the removal of duplicated material could shorten the manuscript overall.

- *We have re-read the manuscript and shortened individual sentences and parts of paragraphs, as appropriate. We would like to note that we disagree with the tone of the reviewer's comment that appears to suggest a lot of repetition and duplication, whilst clearly all sections and subsections are distinct. The listing of advantages in the last part of the discussion is kept on purpose as it closes the contextual framing of open challenges in the discussion.*

2. Some symbols and notations are introduced before being defined.

- *Yes, thanks, we have corrected this.*

3. Variable naming is not always consistent, e.g. mi is used to define the material carried by each simulation particle in equation 3, but pathogen viability in equations 11 and 12, and mv is used to define pathogen viability decay in equation 1, and sometimes m and M are used to define particle mass.

- *We disagree, the naming of the variables noted above is consistent: m_i is used throughout for the viable material carried by each simulation particle. In equation 3 it is introduced and in equation 11 and 12 we describe how this material changes as part of viability decay. m_v is used to denote in conceptual terms a pathogen viability decay function. M is used to denote the total material per source.*
- *We have introduced a table with key parameters to improve clarity regarding notation.*

4. Tables 2 and 3 are information- Highlighting key numbers (e.g. in bold) would help.

- *We had already included bold-faced fonts for selected numbers in tables 2 and 3. To further improve legibility, we have included color shading for the main column.*

5. L22: The sentence "We show that AgPaDS maintains good agreement with HYSPLIT" is overstating validation, suggest changing to something like "We show that AgPaDS produces atmospheric transport patterns broadly consistent with HYSPLIT across a set of benchmark cases"

- *Adapted to more careful phrasing, noting that we do not feel that the previous statement was claiming validation but rather summarizing the comparison to another model.*

6. L697-698: The sentence "moderately longer mean distances from the source (24%), compared with HYSPLIT" should qualify the implications, e.g. "which may be consequential for applications sensitive to arrival timing or long-range thresholds"

- *Agreed and included, but in discussion rather than description of results.*

7. The figure on page 40 of the supplement is incorrectly labelled Figure S3 instead of S35.

- *Corrected, thanks*

8. The figure on page 32 of the supplement is incorrectly labelled Figure S272 instead of S27.

- *Corrected, thanks*

9. Comparison is misspelt in the caption for Figure S29 in the supplement.

- *Corrected, thanks*

10. Marocco should be Morocco in the supplement

- *Corrected, thanks*

11. The caption for Figure S3 in the supplement contains duplicate "(ii)".

- *Corrected, thanks*

12. Many of the figures in the supplement are very small and the text is difficult to read

- *Maximized as much as possible by increasing figure size, noting that we deem it important to have the different inlays/maps next to each other in individual figures per source, and we provide sufficient quality for zoom.*

13. The supplement table of contents lists figures and tables by their page numbers, but no page numbers are provided in the supplement.

- *Corrected, thanks.*

10 May 2026

Response to RC1 & RC 2 – GMD 2026-429

**Citation**: https://doi.org/10.5194/egusphere-2026-429-RC2