10 May 2026

Response to RC 2 – GMD 2026-429

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 RC 2 – GMD 2026-429

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