

This manuscript presents BinMod1d, a package for atmospheric coagulation and breakup using spectral methods. The paper is clearly written and the figures are intuitive. Overall, this is a useful software contribution that is appropriate for the target journal, filling an important niche for an accessible, open-source spectral-bin microphysics model implemented in a modern computing language. The manuscript provides relevant use cases, validation experiments, and demonstrations of the code's capabilities, and the software should be of interest to a broad community of researchers working in cloud microphysics and precipitation modeling.

The instructions for setting up and running the simulations were clear, and configuring the computational environment was straightforward. The figures were reproduced without difficulty and were consistent with those presented in the paper.

I recommend publication after the following suggestions are addressed.

Thank you very much for your comments and suggestions. I also want to thank you for taking the time to install and verify BinMod1D results in your computing environment. I have addressed your comments below in blue text.

Specific Comments

The parameter `sbin` appears throughout the manuscript but is not clearly defined upon first use. Readers unfamiliar with the author's notation must infer its meaning from the equation in line 197 and the subsequent discussion. Please provide an explicit definition (e.g., "number of bins per mass doubling") when first introduced.

I have rewritten this section by explicitly making the mass-doubling grid formula a separate equation with an explicit description of `sbin` as the number of bins per mass doubling. Note, that I have replaced all mentions of "s" with " s_{bin} " for consistency.

L183: More detail on the parallelization strategy would be helpful. Since multiple collision pairs may access the same source and destination bins, it is not immediately clear how race conditions are avoided. Are collision rates evaluated from a frozen particle distribution at the beginning of the timestep, or are bin populations updated during the collision loop? A brief description of the strategy would help readers assess the scalability of the implementation.

To answer your question: it's the first approach for coalescence and the second approach for breakup. BinMod1D calculates the changes in bin mass/number for each bin-pair interaction as *flattened* arrays for all bins (and heights). For coalescence, the transfer is performed via *scatter adding* using the numpy [`add.at`](#) method and the recorded bin pair index for the bin-pair. For breakup, however, an additional Numba loop is used where the gains are calculated for *each thread* such that arrays are shaped (`n_threads` x heights x bins). Another loop then is used to do the transfer after to reduce the (`n_threads` x heights x bins) to (heights x bins) for each category. This is because the total breakup gain terms are

calculated from the loss terms that's calculated in the coalescence section. From trial and error, I found that this strategy was faster than the other methods I tried. I have incorporated this information in the manuscript as follows. Note that I reorganized this section to first describe the Wang et al. numerical integration methodology and then to describe the parallelization strategy at the end of section 2.1:

"To simplify the calculations, each I_k^r term is computed *separately* and then recombined for each k bin for each distribution. This is because each i-j interaction will generate gains in multiple, scattered bins. To perform these calculations, BinMod1D first flattens all arrays needed for each bin-pair interaction such that each bin-pair interaction has a *unique* set of integration values and destination bin indices. Then, BinMod1D efficiently utilizes "just-in-time" (JIT) compilation using the Numba Python package to perform the interaction integration calculations in *parallel for all bin-pair interactions at all heights at once*. This leads to BinMod1D effectively using machine code to loop through each possible bin-pair interaction which, in turn, leads to fast calculations with very low memory usage even when utilizing multiple distributions in the full 1D column model mode. For coalescence, the loss and gain transfer calculations are performed for each bin-pair interaction first using Numba-compiled loops and are stored in arrays with the length of the total number of bin pairs. **The transfer of the calculation from the i-j source space to the k and k+1 target space (Figure 2) is performed through scattered addition using the Numpy add.at method after all integration calculations are performed and stored. Collisional breakup, on the other hand, uses additional Numba-compiled loops after the breakup gain amounts are calculated to transfer the breakup fraction of the total i-j bin-pair source loss to each k target bin (see Figure 3) in the breakup distribution. Multidimensional temporary arrays are used to prevent race conditions during this transfer process where the Numba thread number, heights, and bins represent a three dimensional buffer array from which the transfer sequentially adds to each bin and height in the fragment distribution index for M_k (and N_k if 2 moments are used).** Several core BinMod1D components, including the parallelizable interaction kernel computations, were algorithmically streamlined using Google Gemini Pro as a code-refinement tool to ensure maximum computational speed in standard Python environments."

L316–324: The manuscript describes the use of multiple particle categories and destination-category routing through "cc_dest" and "br_dest". However, I think the paper would benefit from expanding on this, as well as an earlier description of how the kernels or efficiencies may change between classes. Please clarify whether kernels (or probably more importantly the associated efficiencies) are evaluated specifically for each category pair, and if all category pairs are allowed to interact. The flexibility is evident, but I think some recommendations for specifying these interactions would also be helpful, or at least a discussion of the default behavior when category-pair-specific information is not provided.

Thank you for this suggestion. While I wrote BinMod1D v.1.0.10 to only handle a *single kernel type for all categories*, the methodology itself actually utilizes kernel-based arrays with shape (pnum x Hlen x bins) where the zeroth index corresponds to the category pair with a total combination number of $pnum = (dnum * (dnum + 1)) / 2$, and dnum is the number of separate categories/distributions. Therefore, incorporating multiple kernels for each *combination pair* should not be particularly difficult as these kernels can simply be appropriately stacked along the zeroth index. However, I will need to think about how to best implement this change to ensure that users can easily and flexibly use multiple kernels. This change will be important for BinMod1D v.2.0.0 which will have mixed-phase particles and single-particle processes such as condensation, evaporation, sublimation, and melting.

I have included this clarification in the manuscript in section 2.1.

“As can be seen from Equation 4 and Equation 5, all integrals use the same collection kernel $K(x,y|i,j)$ where BinMod1D v1.0.10 considers only a single type of kernel (e.g., hydrodynamic) between each combination of distribution categories.”

Additionally, I’ve expanded upon the cc_dest and br_dest descriptions in section 2.2 (see bold):

“Users can specify the number and type of distributions used in spectral_1d() by adding subsequent distributions to the list in the habit_params input parameter (e.g., habit_params=["snow","fragments"]). The length of habit_params determines the number of separate distributions that are solved where the names refer to the particle microphysical parameters in [habits.py](#) that are used for each distribution. **The spectral_1d() input parameter kernel determines the type of interaction kernel (i.e., $K(x,y)$) with the kernel input (e.g., kernel="Hydro", for the hydrodynamic or gravitational kernel). Users can explicitly set the cc_dest and br_dest input parameters when more than one distribution is specified. Here, cc_dest and br_dest determine the distribution destination of coalesced and breakup bin-pair interactions, respectively where integer values are given from 1 to len(habit_params). For example, a two-category BinMod1D simulation where the first distribution represents a primary category such as ice crystals and the second category represents snow aggregates would be initialized with cc_dest=2 and habit_params=["snow","snow"]. Similarly, a two-category BinMod1D simulation with snow (and its aggregates) as one category and breakup fragments as the second category can be specified with cc_dest=1, br_dest=2, and habit_params = ["snow", "fragments"]. Only one distribution/category is used when cc_dest= br_dest=1 (the default behavior).** Collision (Ecol), coalescence (Es) and breakup (Eb) efficiencies are treated as constants that can be passed into spectral_1d() directly (e.g., spectral_1d(Ecol=1.0,Es=0.5,Eb=0.5)). Here, the total collision-coalescence efficiency is given by $E_{CC} = E_{col} E_s$ whereas the total breakup efficiency E_{BR} is determined by: $E_{BR} = E_{col} (1-E_s) E_b$. In this way, E_b represents the fraction of collisions that occur that do not coalesce but do break up.

Therefore, $E_b=1.0$ indicates the impossibility of rebounding whereas $0 < E_b < 1$ indicates a combination of breakup and rebounding.”

Eq. 18 Specify if D_x needs to be the larger size.

In the revised version, I've specified under equation 18 (now equation 19) that $D_x \geq D_y$. While technically, the Long kernel *does* make the distinction between the larger and smaller sizes, this restriction is not currently incorporated in the code. However, this distinction doesn't matter for self-collection. Therefore, the results presented in the manuscript (i.e., Section 4, Figures 11 and 12, and Table 1) are still valid. Future updates to BinMod1D will ensure that the Long kernel has this distinction for scientific consistency with the Long kernel.

L106: Since the manuscript highlights the advantages of a Python implementation and this paragraph details the recent trend towards accessible languages, I suggest including citations of other Python-based microphysics frameworks that have coagulation, including Particula (<https://github.com/Gorkowski/particula>) and PySDM (<https://github.com/open-atmos/PySDM>). While these models use a Lagrangian rather than spectral-bin approach, they provide useful context for readers interested in open-source Python tools for cloud microphysics. It is also worth noting the language of LCM1d in L110.

I have included the Particula and PySDM references in this section as well. I have also specified the programming language as Python for each package/model as well.

While the repository includes tests for initialization, conservation, and basic diagnostics, the automated test suite appears to contain relatively few tests tied directly to the scientific validation presented in the manuscript. Given the emphasis on analytical benchmark solutions, breakup parameterizations, radar diagnostics, and multi-category interactions, I encourage the authors to consider incorporating more tests based on these published validation cases.

Thank you for this suggestion. The tests/ directory is meant to be strictly unit-based tests to ensure proper installation of BinMod1D and not scientific-oriented tests. The Jupyter notebook examples, on the other hand, are intended to provide these sorts of tests and are also meant to introduce users to the capabilities and use-cases of BinMod1D. I will definitely incorporate more scientific validation type tests in the second version of BinMod1D which will have single particle processes and other features not currently present in version 1.0.10.

Summary

This paper presents BinMod1D, a Python package for simulating collisional coalescence and breakup processes in precipitation, with corresponding polarimetric radar signatures. The model can run as a box model, 1D steady-state, or full time/height column model.

Overall Assessment: The paper is well-written and presents valuable tools for the cloud microphysics and radar meteorology community. The work is solid and the implementation is useful. Future work demonstrating specific applications and case studies will help the community better understand where this modeling approach fits within the broader landscape of microphysical modeling tools and how it can generate the most scientific impact.

Thank you very much for your comments and suggestions. I have addressed your comments below in blue text.

Comments

General Comment on Figures 1-3: The text describing the methodology is quite dense, and while the figure annotations are good, the narrative could be expanded to better guide readers through interpreting these complex diagrams. Consider adding:

- Step-by-step walkthrough of what each panel shows
- More explicit connections between figure elements and the mathematical terms in equations
- A "how to read this figure" text, especially for Figure 1's complex gain/loss regions
- Clearer transitions from figure to text and back

The figures contain important conceptual information about the numerical scheme, but readers may struggle to extract insights without more narrative scaffolding.

Thank you for these suggestions. I have added more detail to the Figure 1-3 narrative and equations to better explain these figures and equations and how these methodologies work. I've also slightly modified Figure 1a to better illustrate how the source bins correspond to the gain region. Some of your suggestions would be difficult to directly implement. For example, it is difficult to provide step-by-step walkthroughs as the figures themselves represent a *single step* in the numerical integration methodology. However, I have added more text to better explain these figures. In particular, I added the following components:

1. I modified Figure 1 to include hashed source (i) bins for the target k bin to highlight that both i and j bins are being included in the k bin gain.
2. In the text describing Figure 1, I incorporated some of the mathematical details regarding how the x+y line relates to the x_k and x_{k+1} k bin edges as well as how this creates the red “Gain” region.
3. I added parenthetical explanations in equation 9a-d which describe how each integration term corresponds to number/mass gain and loss transfer amounts.

Discussion

The paper presents a solid technical contribution, but would benefit from expanded discussion of:

- Where BinMod1D fits within the existing ecosystem of microphysical modeling tools
- Specific use cases where this approach has advantages over alternatives.
- Potential applications in operational radar retrieval development or research workflows
- Future work demonstrating community impact through real-world case studies

This would help readers understand not just **what** the tool does, but **when** and **why** they should use it.

Thank you for your suggestions. I’ve added additional text to the conclusions section which touches upon these points. I would like to point out that the original version does already have some of these elements. For example, in the current preprint version, I do mention the possibility of using BinMod1D for investigating the evolution of snow size distributions using initial PSDs taken from in situ aircraft or laboratory measurements (similar to the Field et al. 2006 study and Connolly et al. 2012 study, for example). In the conclusions section of the current preprint, I also provide a quite detailed BinMod1D example which describes a potential application of radar retrieval development as well. However, I have added a few additional sentences in the conclusions that address these suggestions:

“BinMod1D is a simple, economic spectral bin microphysics model that should be accessible for those familiar with the Python programming language. **Other similar Python packages like LCM1D and PySDM use Lagrangian depictions of the PSD evolution rather than the typical Eulerian description given by traditional spectral bin microphysics schemes. In this way, BinMod1D represents a much more traditional approach that is more similar to the widely-used Fast Spectral Bin Microphysics scheme (FSBM; Khain and Lynn 2009) which is available in the Fortran-based Weather Research and Forecasting (WRF) model. BinMod1D provides an expedient way to not only generate box model or 1D column model PSD evolution due to**

coalescence and/or breakup, but also to easily incorporate various habit parameters, numbers of distributions, and interaction kernels. Additionally, BinMod1D, unlike other similar Python packages, allows for direct, explicit simulation of steady-state vertical radar profiles. This article provides only a small set of examples for using the model. However, due to its flexibility, BinMod1D can be customized for a variety of different purposes. For example, users could use BinMod1D to perform in-situ-type bin model simulations similar to Field et al. (2006) where actual in-situ-derived particle size distributions are used to initialize the bin model and investigate the theoretical effects of aggregation (and/or breakup) compared to measured distributions taken from such Lagrangian aircraft spirals. Similarly, users could also develop theoretical particle size distributions valid for laboratory studies such as done in Connolly et al. (2012)."

Minor Comments

- Line 127: Missing text - should read "model, due to" (I think)

From the online preprint, I believe the line in question is this:

"Therefore, the ability of BinMod1D to directly simulate steady-state vertical profiles drastically decreases the computational time compared to running the 1D column model to its steady state as done in several rainshaft studies (e.g., Prat and Barros, 2007b; Prat et al., 2012; Kumjian and Prat, 2014)."

If I'm looking at the correct line here, I believe what I have written is correct. The point here is that BinMod1D allows users to directly calculate steady-state (in height) profiles rather than running a time/height column model out to its steady-state solution (using a sufficiently long time period). Please let me know if this is not the line in question or if there was a different concern here.
