

Reviewer 1:

We sincerely thank the reviewer for taking the time to thoroughly review our manuscript. The detailed feedback and constructive suggestions are very appreciated. We believe to have made the necessary revisions to address all comments.

In the following, responses are in blue, and quoted text is shown in green. Text after the little arrow '→' is newly introduced or modified manuscript text in the reaction to the reviewer's comments.

General comments

This paper introduces LARA (Lagrangian Reanalysis), a novel global dataset created by converting ERA5 reanalysis data into a Lagrangian format using the FLEXPART model. Unlike traditional Eulerian reanalyses, LARA tracks six million air particles from 1940 to 2023, providing high-resolution, hourly data on their positions and atmospheric properties. This dataset supports studies on atmospheric transport, energy and moisture fluxes, and extreme events. It is validated for internal consistency. LARA is available in Zarr data format, accompanied by analysis scripts and four illustrative applications, such as Hadley cell evolution and diagnostics of data assimilation discontinuities in ERA5. Overall, LARA appears to be a powerful tool for Lagrangian-based atmospheric research.

The methodology of the study is sound. FLEXPART is a reliable model, and the transition to a Lagrangian perspective is well-justified. The validation is appropriate, and the analysis is carefully executed, offering a valuable dataset with wide applications.

One general comment that remains open is how the number of six million air parcels used in LARA was determined. This number is briefly discussed in the conclusion but seems based on computational and storage constraints rather than a sensitivity analysis? Clarifying whether tests were conducted to assess the adequacy of this particle count would strengthen the manuscript. It would also be useful to discuss whether scientific results in Section 4 would remain robust with different particle counts.

It is indeed constrained by mainly the storage. The number of particles sets a lower limit on the volume in time and space needed to sample enough particles to obtain statistically significant results. This is described in Section 3.2. Since all our examples use large volumes, these results will not be affected by increasing the particle number, although it would make some of the maps shown smoother and possibly reduce some noise. We added the following statement to Section 2.3:

The dataset is large, with a total size of approximately 320 TB, and making it openly available constrained our choice of total particle number and temporal output resolution. Section 3.2 describes possible limitations on the use of the dataset due to its resolution.

Overall, the manuscript is well written, scientifically rigorous, and uses proper terminology and referencing. It follows a clear structure and presents concepts and results in an accessible and precise manner. While some sections are detailed, this aligns with the paper's dual role as a dataset description and usage guide. Overall, the manuscript meets

high standards of clarity, conciseness, and scientific communication. I recommend publication subject to addressing the specific comments and technical corrections listed below.

Specific comments

Lines 34-39: Several major reanalyses, including MERRA and JRA-55, are mentioned, but it would be useful to note that newer versions are available, MERRA-2 and JRA-3Q. Citing a recent reanalysis intercomparison, such as from the S-RIP/A-RIP initiatives, would provide broader context on how these datasets compare in performance and scope.

The citations and mention of MERRA-2, JRA-3Q, and the intercomparison project are added.

Lines 70-73: The Lagrangian datasets by Sodemann et al. and Vázquez et al. are described as ‘available on request’, but it’s unclear how they can be accessed. If they are not openly available in a FAIR-compliant manner, it may be better to omit or clarify these statements to avoid misleading readers, especially compared to the open-access nature of LARA.

We omitted the statements in the revised version of the paper.

Line 74: The phrase ‘showing large decrease in particle distribution degradation in the troposphere over time’ is somewhat unclear. It would help to briefly explain what is meant by ‘particle distribution degradation’, such as whether it refers to deviations in air density, loss of particle ensemble representativeness, or numerical diffusion effects.

We added the following statement: Since then, a new version of FLEXPART (version 11) has been released, showing a large decrease in particle distribution degradation in the troposphere over time, where in version 10.4, the particle number density does not well represent air density after some time (Bakels et al. 2024).

Line 78: The manuscript refers to the ‘full period of 1940–2024 ERA5 reanalysis’, but the title states 1940–2023. This discrepancy should be addressed for consistency, clarifying if the dataset includes data through March 2024, as mentioned later in the paper.

Thank you for pointing this out. We corrected the time period running until March 2024. Since we only cover 3 months of 2024, we decided to leave the title intact.

Line 82: Instead of providing just a web link, include a proper citation with a DOI for the LARA dataset to ensure long-term accessibility and facilitate referencing, in line with best practices for data publication. At the time of this review, the web link was not accessible, unfortunately. The web site showed ‘Internal Server Error’.

Our apologies, we replaced it with the proper persistent identifier, which also includes a description of the dataset (<https://phaidra.univie.ac.at/o:2121554>).

Table 1: The additional notes in the ‘unit’ column are unclear. For example, it’s not obvious what is meant by etadot having units of s^{-1} but ‘internally: $Pa\ s^{-1}$ ’. If these notes are crucial, move them to footnotes beneath the table for better clarity without cluttering the unit column.

Thank you for the suggestion. The notes have been moved to footnotes.

Line 93: The manuscript mentions using ERA5 data at $0.5^\circ \times 0.5^\circ$ resolution, although ERA5 is (natively) available at $0.25^\circ \times 0.25^\circ$. Clarifying why the reduced resolution was chosen (e.g., due to computational or storage constraints) would help improve transparency regarding the dataset's accuracy.

“For the creation of the LARA dataset, we used the most recent reanalysis dataset of ECMWF, ERA5, with hourly 0.5×0.5 due to storage constraints data as input on 137 vertical model levels, of which 88 are located below 20 km.” → “..., with its native hourly temporal and 137-level vertical resolution and a reduced horizontal resolution of 0.5×0.5 due to storage constraints.”

Line 121: A brief explanation of the ‘domain-filling option’ would help readers unfamiliar with the concept. Clarifying that it initializes particles to uniformly represent the entire atmospheric mass and discussing its relevance to mixing and atmospheric transport would enhance clarity.

The following clarification has been added: “..., which discretises the full atmosphere into particles of equal mass. The particles are initialized such that their number density is proportional to the density of air, and each of them represents an equal fraction of the total atmospheric mass.”

Lines 126–128: Since the data processing was split into overlapping streams, clarify if consistency checks were done across stream boundaries. It would be useful to mention whether users need to account for potential artifacts or discontinuities when working across different stream periods.

Our apologies, there were some mistakes in the text that are now clarified. There is a full year overlap, since particles cannot be traced between the different streams. The full year of overlap guarantees that particles can always be traced - also across the boundaries - below the span of one year (corrected in the text), which is more than enough for most applications. The difference in ‘quality’ of the overlapping years at the start and the end of the stream is shown in section 3.1.

Lines 126–130: The explanation of overlapping periods between computational streams is unclear. The text first states ‘one full year for each period’, then mentions a ‘three-month overlap’, which is confusing. Please clarify whether the overlap is one year or three months, and explain how it ensures continuity in trajectory tracking.

Apologies, this is corrected and clarified now, see above.

Lines 132–136: While the impact of time step length on trajectory accuracy is discussed, the numerical integration scheme used in FLEXPART is not mentioned. Specifying the integration method (e.g., Runge-Kutta, Euler) would provide a more complete understanding of the model setup and accuracy.

The following statement is added to the FLEXPART section 2.2: “FLEXPART uses a first order Euler method in addition to a Petterssen correction (Petterssen, 1940) for advancing particles.”

Lines 139–141: The inclusion of compiler flags and OpenMP settings may be overly detailed for an ESSD paper. Unless these settings impact the results or reproducibility, consider omitting it for clarity and focus.

The flags have been omitted.

Lines 144–145: The dataset provides hourly means for several variables, but adding additional statistics like minimum, maximum, or standard deviation could be valuable, especially for applications involving extremes or uncertainty. If not included, note whether these could be added in future versions.

We considered adding such values, but due to storage constraints decided against it. However, some variability could be derived from comparing the instantaneous values to the averaged values. Uncertainties could be derived from comparing groups of particles to each other given the same volume, and maxima and minima should emerge when using large enough numbers of instantaneous particle values.

Lines 157–161: Include an estimate of the total size (e.g., in terabytes) of the full LARA dataset, especially since storage format and compression are discussed. This would help users understand data management requirements for downloading or processing large subsets.

A statement has been added (see first comment).

Line 170: Perhaps briefly elaborate on improvements to the FLEXPART interpolation scheme in the latest version. Summarize enhancements to help readers understand their role in reducing potential errors.

We added the following clarification: “Using the improved interpolation scheme of FLEXPART 11, a reduction of ~15% and ~8% in PV errors is observed above 10 km and between 5-10 km, respectively.” → “Using the improved interpolation scheme of FLEXPART 11, which removed an interpolation step by advancing particles using the native vertical coordinate grid of the ERA5 data, a reduction of ~15% and ~8% in PV errors is observed above 10 km and between 5-10 km, respectively.” We also added: “More details on FLEXPART 11 improvements can be found in Bakels et al. (2024).”

Figure 1: The x-axis label refers to absolute differences, $|\rho_{\text{part}} - \rho_{\text{ERA5}}|$, but negative values suggest signed differences ($\rho_{\text{part}} - \rho_{\text{ERA5}}$) are shown. Update the label to reflect this. Also, suggest to change the y-axis label to ‘Geopotential height (m)’ for clarity.

This has been corrected, thank you.

Lines 193-194: The sentence about maximum deviation of 0.6% from latitudinal-averaged volume is unclear. Clarify what is meant by ‘latitudinal-averaged volume’ and how the 0.6% deviation is calculated.

This is a mistake, it is the deviation over the whole latitudinal volume, with longitudes spaced by 0.5 degrees. We decided to remove the paragraph: “The maximum deviation ... can be larger.” since it does not provide any new information and potentially causes confusion instead.

Line 199: Suggest to rephrase ‘it is always better to take the air density from the _ERA5_ values interpolated to the particle positions’ for clarity.

Rephrased as suggested.

Lines 201-202: The sentence about calculating mass using average air density within the source or receptor volume is unclear. Does this affect mass conservation? Should users always rely on air density or mass derived from ERA5 data at the particle locations, rather than using the mass of Lagrangian particles?

We removed this sentence and the next (“Similarly, when the total transported mass ... source or receptor volume.”), since they cause more confusion than they are helpful. Figure 1 should be sufficient for explaining the particle density discrepancy.

Figure 2: The figure seems somewhat overloaded with information. Consider splitting it into two figures for panels a and b or removing unnecessary curves/data to improve clarity.

The figure has been simplified, removing unnecessary curves, while also putting more emphasis on explaining the methods. We separated the figure into three steps: 1) maps of the selected particles; 2) Particle distribution and PDFs; and 3) an analysis that could be conducted using these distributions. We hope this makes the figure more valuable in a pedagogical way.

Lines 263-266 and 273-276: The phrasing could be adjusted to avoid casting doubt on previous results. If only local/regional analysis is meaningful, why focus on global results in Fig. 2? A more confident presentation of the findings could strengthen this section.

Thank you for pointing this out. We have rephrased the paragraph in a more positive light, explaining why we do not go in much detail for the examples: “To demonstrate the potential of LARA, we only present global results, but it would be interesting to extend the analysis to particular regions (i.e. Pacific, Atlantic, Indian Ocean). This would then allow for the investigation of other processes, e.g. global warming, ...”.

Lines 299-300: Could convection parametrizations influence the results?

*This is unlikely. Eckhardt et al. (2004) did not use a convection scheme for their trajectory calculations, while here we did, but the results are very similar. The text read as if convection parameterisation was applied to Eckhardt et al. (2004), which is incorrect, and we therefore added the following clarification to the text: “**Eckhardt et al. (2004) did not use a model that parameterises convection, either. However, when combined with the requirement that the air mass within a WCB be greater than 1 kg per square metre - which implies a certain degree of coherence in airstream motion - it reasonably corrects for the erroneous classification of deep convection.**”*

Lines 318-320: Why use a 1000 m proxy for boundary layer height, given that the actual boundary layer height is available in the LARA dataset?

We agree it might have been better to use the actual boundary layer height. A more in-depth follow-on analysis (Brack et al., in preparation) actually used boundary layer heights. For the purpose of demonstration, however, we consider the use of a constant height sufficient.

Figure 5: The tracer conservation errors are averaged up to 20 km altitude, but tracers considered here vary exponentially with height, whereas particle numbers decrease exponentially with height. Is there any significant altitude dependence in tracer conservation errors? It could be insightful to compare conservation errors across different layers (e.g., lower, middle, upper troposphere, and lower stratosphere).

This would indeed be very interesting. As earlier shown in the FLEXPART version 11 paper (Bakels et al. 2024), the altitude makes a large difference in conservation error. However, our aim for this paper is to show possible avenues of what could be done with the dataset, without going into much detail. Of course, we hope these questions will be picked up by others and be further explored upon.

Line 367: Add '(not shown)' since this is not demonstrated in Fig. 5.

Added “(not shown)”.

Lines 377-378: The improvement in consistency may be due to the broader onset of the modern satellite era, not just the TOVS instrument.

We improved the statement including the broader onset of the modern satellite era by replacing the following: “This is in line with the introduction of assimilation of data from the TOVS (TIROS Operational Vertical Sounder) satellite series at the end of 1978.” → “This is in line with the introduction of assimilation of satellite data, such as from the TOVS (TIROS Operational Vertical Sounder) satellite series at the end of 1978.”

Technical corrections

Please ensure consistent capitalization and abbreviations for terms like 'Section' and 'Figure' throughout the text, following Copernicus manuscript guidelines.

Thank you, this has now been corrected.

Line 69: rephrase 'ERAInterim' as 'ERA-Interim'

Corrected.

Line 83: please check and standardize the spelling of 'Zarr' data format throughout the manuscript

It has been standardised now.

Line 133: 'deviate by approximately 20% less from ERA5 air densities' – Do you mean '20% or less'?

We clarified it by replacing the sentence as follows: “We found that particle densities deviate approximately 20% less from ERA5 air densities when using 300 seconds time steps compared to 600 second time steps.”→ “We found that particle densities show smaller deviations, with a reduction of the difference of particle mass densities to ERA5 air densities of approximately 20% when using 300 seconds time steps compared to 600 second time steps.”

Line 141: Do you mean '12 eight-year periods' (considering the overlap)?

Thank you, we corrected the mistake.

Table 2: Column header 'unit' -> 'Unit'. Fix unit 'K' for tropopause height.

Thank you, it is corrected.

Line 254: The acronym 'ERSSTV5' is not introduced.

We added the full name (Extended Reconstructed Sea Surface Temperature version 5).

Line 285: model used in _that_ study

Corrected

Figure 5: suggest to replace '-23°N' by '23°S' in the caption

Corrected

Lines 382-383: Please check the capitalization of 'Northern Hemisphere' and 'Southern Hemisphere' throughout the manuscript for consistency.

We have now consistently capitalised it.