Dear Editor and Reviewers,

we thank you for the time and effort you invested in handling and reviewing our manuscript, "Seismo-Acoustic Meteoroid Observation Recording Database (SMORD): A Global Dataset and Deep-Learning Phase Picker for Meteoroid-Generated Air-to-Ground Coupled Seismic Waves". We appreciate the constructive feedback and thoughtful suggestions, which have helped us to improve the clarity, and overall quality of the paper.

We have revised the manuscript accordingly and provide a detailed, point-by-point response below. For each comment, we first restate the reviewer's remark and then describe our response and the corresponding changes made in the revised manuscript (highlighted in light blue).

We hope that the revisions adequately address the concerns raised and that the manuscript is now suitable for publication Seismica.

Sincerely,
Dario Eickhoff, Runa Ostermeier, Joachim Ritter

Editor (Kenneth Macpherson):

I had a small additional request. In line 188 you mention interpolating to a common sampling rate of 100 sps, and on line 151 you mention bandpasses up to 40 Hz. For many seismic stations, including the TA example in the manuscript, this would mean up-sampling and filtering above the Nyquist of the original sampling rate. Can you explain why the 100 sps is necessary, and how you limit artifacts that may be introduced by filtering beyond the Nyquist?

The common sampling to 100 Hz is used for the training because PhaseNet expects 100 Hz sampled data. However, as training is done on raw unfiltered waveforms the bandpass used during the manual phase identification never affects the training data. We added a clarification in the "Model training" subsection to clarify that the training and resampling happens on unprocessed waveforms.

Reviewer 1 (Christoph Pilger):

Dear author and editor,

the manuscript about SMORD by Eickhoff et al. provides description, application and discussion of a phase picking method for meteoroid signal detection in seismometer data using deep learning algorithms based on a training of manually identified events. The study is well written and a relevant investigation providing an important step forward from manual to automatic identification of meteoroid signals in seismic detections. It is suitable for publication in Seismica following revision with remarks provided below.

The abstract and non-technical summary well describe the study's content and results. The introduction provides largely sufficient context for meteoroid entries into Earth's atmosphere, related acoustic shock waves and acoustic-to-seismic coupling, preparing the ground for the actual study. The data and methods sections provide the necessary basis for the given investigation, including suitable referencing (it can be extended at some minor points, as described below). The result and application sections contain well selected figures of good quality and provide a reasonable overview of the database, outcome of the investigation and exemplary use case (with some room for improvements as commented below). The discussion and conclusions are well elaborated and put SMORD and its potential benefit in a larger context, providing numerous potential future applications and ongoing work to get there. The final statement about data availability upon publication of this study is very promising.

Detailed comments and corrections for further improvement of the manuscript are:

L67-74: “are studied by different techniques”: I would add infrasound measurements here for completeness as well. I saw they are described later on, but they should be mentioned here as well for the full picture. Suggestions for references: Ens et al., 2012 (https://doi.org/10.1016/j.jastp.2012.01.018), Ott et al, 2019 (https://doi.org/10.1016/j.pss.2019.104715) or Silber 2024 (https://doi.org/10.3390/rs16193628).

We agree and added infrasound to the techniques and added appropriate references as requested.

L80: “seismic waves are excited with much higher propagation velocities.” ...higher than sound velocities, yes, but since you also speak of the 11-72 km/s, which is faster than seismic velocities, please rephrase/clarify the sentence (or quantify the seismic velocities).

Done.

L82 (Figure 1): can you add the lateral distance to the caption or subfigure 1c to get a general impression of the ranges for this case (either range from xx km to xx km to the caption or some additional y-axis labeling).

We added the lateral distance information to the subfigure 1c and reordered the seismic waveforms to align with increasing distance in y-axis direction.

L93: “record seismic waves”: this implies to a certain degree that the coupling takes place remotely and the meteoroid signal propagates as seismic wave (of which type, the reader might ask then, surface wave?). Please clarify, is this observed or meant? Or is it coupling to the ground directly at the seismometer location (L161 “plausible moveout of an acoustic wave” supports this suggestion)... then “ground motion” might be a better term here and some explanation of the coupling process might help. Fernando and Charalambous (2026, https://doi.org/10.1126/science.adz4676) call it in a recent article (for space debris reentry signals): “impulsive downward peaks in the ground velocity were recorded. These were N waves, or near-instantaneous downward compression of the ground in response to the incident sonic boom, followed by upward rebound during subsequent rarefaction”. You also gave and referenced a description of the coupling, but much later in L448-450, this could also be moved or repeated already here.

We clarified the paragraph on the recorded seismic waves and added some of the references from the Discussion.

L103: “1 kg”: please harmonize with L92 “10^3 g”. Just to clarify, you could mention (L92 maybe) that this is the mass when entering the atmosphere.

Done.

L107/108: “In addition, a rapid calculation of the trajectory can help to recover the extraterrestrial material for further studies”: since this is the first time in the introduction that you mention the trajectory, it would be good to elaborate in some more detail, what has been done and what can be done regarding trajectory estimation (by available observations in general, and using seismic data in this case), adding references.

Done.

L136/137: “includes re-entries of space debris, which produce seismo-acoustic signals analogous to meteoroids”: citation recommended, e.g. Fernando and Charalambous, 2026 (https://doi.org/10.1126/science.adz4676) or Fernando et al., 2024 (https://doi.org/10.26443/seismica.v3i1.1154). How many of the cases are flagged as space debris?

We added the number of included reentries (9) and the references.

L138: “CNEOS fireball catalog”: analog to IMO I recommend a short description, which observation sources are used for the CNEOS catalogue, if possible with a citation as well.

Done.

L146-148: some brackets, especially “(“ are missing, or some “)” are too much.

Done.

L154: “2. Methods” -> “3. Methods”

Done.

L156: “around expected wave arrivals” ... is this as first guess a wide time window considering anything coming in around/after the expected origin time? Or is it already more specific taking into account certain (e.g. acoustic) travel times? Can you quantify these expectations in minutes/hours, or is it a very individual case-by-case situation?

We added clarifications to explain the way the time windows are set.

L171 (Table 2): Table entries “P phase arrival sample” and “Split” are unclear without further explanation.

We added explanations and a reference to the following section to clarify the table entries.

L223 “increase the amount of training data for each epoch”: can you quantify the amount of training data (as you did with the number of events, labels and the percentage of training cases), i.e. how many variations using rotation and noise are created?

We added the amount of data for one and all three generators.

L262 “across a global station set”: is it possible to name or estimate the number of stations taken into account, or is this time-dependently very dynamic?

We added the number of unique stations to the text.

L269 (Figure 2 caption): Please clarify what the 2025 means in brackets after IMO and CNEOS. The color coding goes until 2024, so is it the retrieval date? If it is a citation, this is not clear, since not in the reference list like this.

We corrected the references to the IMO and CNEOS catalogs.

L291-297: About the Perseids. Probably the necessary 100 IMO sightings are not reported for a single event during a meteor shower, because observers see (and expect to see) so many meteors then. This is what you also notice in your manuscript. But can it also be that the meteoroid events strong enough to generate signals at seismometers are more singular events and naturally disconnected from any meteor shower activity? Do you have indications that a number of the events you have in your database are connected to other meteor showers (or not)?

We added an explanation that currently no meteor shower association is done.

L310-314: Can you spend a few words, if the signals you detect are completely emitted from a generally closest point of the trajectory rectangular to your station, or if specific fragmentation points are the source(s) of your seismic signals.

We added some clarifications to the possible meteoroid sources.

L346 (Figures 4c and d): I do not fully understand the histogram(s) in figure 4c: Why do I see one big bar up to 270 and in the inlay figure a bar up to 200 (so this is not a zoom of the total figure). The dashed lines are the P_5 and P_95 values, but please better explain the bars, bin size and y-axis value of the histogram(s). Figure 4d: the blue and red curves are normalized probabilities? Is the value of “true” probability of 1 shown by the corresponding blue dashed lines, or do these just separate subset of lines and have no other meaning? Please modify the caption (and/or text) to better describe it.

We added an explanation to subplot 4c. It shows a zoomed in portion with different bar widths to better show the distribution of errors around 0 seconds. We have removed the blue dashed lines.

L353 (Figure 5b): in the right figure y-axis label “signal-to-noie” there is a “s” missing. Regarding the y-axis unit “in dB”, this unit should probably better be explained (20 dB, factor of 10) or labeled differently, since not used or explained anywhere else in the manuscript. And in the color bar I would expect +/- 0.75 instead of +/- 0.65 at the respective ticks, at least it is following linear scaling.

We have corrected the y-axis label, the colorbar ticks and changed the snr from dB to a simple ratio.

L363 or probably to be treated in the discussion section: how much effort is this exemplary application as workflow for each case (L389: "quickly", L479 "rapid")? Could it be a standard (automatic) application and in the end a product to generate a trajectory. Is the non-automatic step 4 always necessary or recommended for this (L384/385: "slight manual curation … can improve") and how much effort is this? How dependent is the method on available CNEOS or FRIPON values? It looks like that the IMO values are less useful and have a much higher error compared to the true trajectory. Since most of your cases are IMO-based only, can the procedure be conducted anyway and can the outcome probably be a benefit (providing a much better trajectory) … probably have a look and take into account for discussion the outcomes of Fernando and Charalambous, 2026 where they also derive a trajectory many kilometers apart from the predictions, using seismometers.

We added a sentence in the discussion that outlines how much of the 4-step workflow can be automated. The manual curation does not take much time. We know that the IMO-based trajectories are not as useful as for example FRIPON trajectories. This is also why we apply a generous search area radius around the trajectory midpoint. The method that we use to finally calculate the trajectory is robust to uncertainties in the geometry and location of the trajectory. There will also be a future publication that discusses the trajectory calculation methodology in much more detail.

L364: add the distance up to which data was downloaded

The distance and the actual midpoint used have been added.

L369: first time and good place for introducing the abbreviation PSO

Done.

L374 "the one without correction": add Figure 6a or 6a,c,e here, or generalize the Figure 6b at the end of the sentence to 6b,d,f … or just mention "Fig. 6" somewhere and keep it with that.

Done.

L376 -378: Please provide some more details here, especially based on the number values of table 4, which case and configuration is best in your six cases and the comparison (mentioned in L361) and why this is so.

We added a paragraph to highlight the best fitting solution.

L379 (Figure 6): in subfigure 6a there is a \tau symbol pictured on the deviating grey trajectory, which should be removed. In the caption: use “gray lines” instead of “gray bars”. Add some verb behind “seismic station distribution”. Change “in respect to” to “with respect to”. Remove the first usage of the word ‘two’ in: “b, d, f) same as a, c, e), but for case 2 with two PN picks manually verified and two corrected PN picks.”. 6g: difficult to see waveforms and bars clearly here, check for better figure quality here. One of the two blue circles is at probability zero, this means you removed an invalid pick?

We cleaned up the figure and the description and added an explanation for the zero probability pick.

L390 (Table 4): Please describe in caption or text what the RMSE values are for the PSO solutions. Is it deviation to reported origin time, and/or to IMO or FRIPON values, or travel time? Furthermore, please provide some more information about the PSO, how it derives the trajectory angles (is it a forward modeling or back projection, or a travel time error minimizing) and especially how it then derives the velocity. This might be a subchapter in your methods section, unless you want to keep this section restricted to the SMORD procedures and not to a SMORD application (then some more description in section 5 is needed there to make it “stand-alone”).

We added a description for the RMSE.

L416: I suggest to extend “atmospheric winds” to “atmospheric temperature, winds and frequency-dependent attenuation”

Done.

L421 “seismo-acoustic”: previous wording was often acoustic-seismic. I suggest to harmonize the word usage, presuming seismo-acoustic is more known and used compared to acoustic-seismic.

Done.

L441/442: “The observational basis, continuous online waveform data,”: I suggest to add a “which is” after the comma

Done.

L445 and L447: “acoustic-seismic moveout etc.”: here I would reduce to “acoustic moveout” or “moveout of an acoustic wave” and also skip the “etc.” in the first instance.

Done.

463-471: uncertainty is an important point but probably not a good place for discussing it at the very end of discussion, after all the outlook and future application paragraphs. Lines 434-439 already treat uncertainties, seems better to move these sentences here.

Done.

L476 “held-out”: what do you mean with this?

This has been changed to the test subset of all events from the SMORD database.

Reviewer 2:

The manuscript submitted by Dario Eickhoff compiles a database of fireball events around the world along with seismic waveforms recording the air-to-ground coupled shock wave produced by the impactor during flight. The database is used to train a machine learning phase detection model that picks time segments out of the waveforms to mark the onset of P-waves excited by the acoustic wave. The authors use an example event to demonstrate how the model can be used jointly with particle swarm optimization to pick the onsets and reconstruct the trajecory of the impactor. The proposed methods could provide a valuable resource to the meteoroid community by providing an accessible seismic record of these events. If developed further, the modified phase picker might be able to automatically register new events and used to quickly estimate the strewn field for meteorite retrieval. Overall, I appreciate the concept and fascinated by the application. The motivation, results, and conclusions appear sound, but missing punctuation, awkward phrasing, and incoherent sentence structure made certain points difficult to understand and distracting. After improving the writing, I think this manuscript could be reconsidered for publication.

-- Notes –

I am not very privy to meteor jargon, but I believe the objects being studied here are technically meteors of the fireball variety, as meteoroids are in space become meteors when they enter the atmosphere.

Meteoroids refer to the rock body if it is still in flight, while meteors are the luminous phenomena produced by ionization of atoms in the atmosphere. Because we measure the acoustic-to-seismic coupled wave produced by the meteoroid rock body we refer to it as that. Optical techniques and observers do observe the meteor and, which are also fireballs if their apparent magnitude exceeds -3.

L78: "72 km s^{-1} (Ceplecha et al., 1998; Moorhead et al., 2019), considerably" added comma use en dash for range of numbers e.g. 290–350 m s^{-1})

Done.

L79: "After coupling into the solid Earth, seismic waves are excited with much higher propagation velocities." Remove the ambiguity. This could be interpreted as the seismic waves excited by the impactor making contact with the ground and not the shock wave as the previous sentence has the fireball as the subject. Higher velocity than the acoustic wave, not the fireball.

This has been corrected.

L82: "Here, seismo-acoustic" consistent panel labels a), b) use en dash: 1–10 Hz

Done.

L83: Split the long sentence and clarify the comparison. Something like: A key advantage of seismo-acoustic observations is the abundance and global distribution of seismic sensors, which provide broad and dense spatial coverage compared with the much smaller number of dedicated meteor cameras. These cameras, in turn, greatly outnumber meteor radar systems.

Done.

L92: comma after 1e3 g

Done.

L95: "lie in the ~1–3% range"

Done.

L96: "site conditions, such as the"

Done.

L100: Condense the following sentences, it's a bit redundant. I think you could get away with saying acoustic waves are detectible from masses greater than 1 kg and seismic waves from masses greater than 1 g. If I understood that correctly. This whole paragraph could benefit from improved flow and concise language.

We have improved the wording and restructured the paragraph slightly.

L107 & 108: extra-terrestrial or extraterrestrial

We chose extraterrestrial.

L109: "For these reasons, we focus on meteoroid events from publicly accessible catalogs [that] contain all significant events" I think you should omit "which contain all significant events". Is not possible there are missing events? I'm not really sure what that is supposed to mean and it is unclear the reasons for choosing these catalogs, because recovering extraterrestrial material and risk assessment?

The “all” has been removed.

L111: consider spelling out CNEOS once here

Done.

L119: "Furthermore, we characterize basic statistics of acoustic-to-seismic coupled waves across a gloabl station set" I don't know what this is referring to, metadata?. I wonder if "we describe the database" covers whatever this is referring to.

We have shortened and removed this sentence from the paragraph.

L123: arrival-time detection and comma after SMORD

“Arrival” was changed to “phase” and the comma has been set.

L135: I never heard "great-circle radius" before. Great-circle distance of 3° (or, within an angular distance of 3°) is more common.

Radius has been changed to distance.

L146: extra right parentheses)

The parentheses are now set correctly.

L154: 3. Methods

Done.

L159: en dash would look nicer

Done.

L168: "... locations, event together with pick times" doesn't make sense. Correct and smooth this sentence out. Table 2: What does a significant and moderate SNR mean? Did you use a threshold like SNR>5 is significant?

Sentence has been smoothed out.

L173: "meteoroids, we trained"

Done.

L180: "we train the network on postive windows centered on analyst labeled arrivals". Don't you also train with negative/noise windows and off-center positive windows?

This has been corrected.

L182: a probabilistic loss function? I'm not sure about that. The output is a probability of arrival, but the probability is calculated deterministically based on the parameters of the neural net. Then the loss function is the difference between the arrival probability and the binary label or the arrival probability after a threshold is applied to give it a binary value and the binary labels, is it not?

This has been corrected. It is a weighted cross-entropy loss that is minimized over the epochs.

L205: centered on labeled arrival/signal onsets. meteoroid onset makes you think of the object itself.

Done.

L225: I don't understand why there is zero padding. Why not just use the original waveform. using zero-padding seems like the signals would have infinite SNR.

Zero padding is solely applied to maintain the fixed input length of 6,000 samples required by the neural network architecture. When the labelled onset is positioned near the boundary of an available waveform, there may be insufficient real samples on one side. Rather than discarding these events, zero padding preserves them during training, thus maintaining dimensional consistency across all input tensors. However, zero padding is a limitation of the current approach, which could be overcome in a future iteration of the model by using real data for padding instead.

L232: "10% of the normalized amplitude for each epoch for one generator" isn't it instead normalized amplitude for each signal? and I don't know what "for one generator" means

We changed the sentence to specifically address generator 2.

L235: For positive examples, we use

Done.

L264: "CNEOS events represent ~17 % of the database and they contribute to ~6 % of labeled picks." How is there only 54 since "all CNEOS events from 1988-2024 with reported estimated yields ≥0.3 kt TNT equivalent" is used. CNEOS events with those magnitudes account for nearly 400 enteries. Is this impact energy or radiated energy? presumably, filtered over land and where seismometers had detections.

We added an explanation for this. Essentially it comes down to the distribution of events, with many events from the CNEOS database happening over the open ocean, where no seismometers are placed nearby.

L266: "a consequence of the eyewitness-driven IMO contribution." I would say seismometer coverage is the dominant factor imposing the bias as there is a hundreds of fireballs around the world from CNEOS and orders of magnitude more people in other places around the world, but most seismometers are in the U.S. and Europe.

While the seismometer coverage certainly plays a role in the spatial bias, the IMO is the larger factor because it constrains the areas where data is looked at in the first place. This is also made apparent in lines 300-302.

L269: "Colored circles mark IMO (2025) events and colored squares mark CNEOS (2025) events." the 2025 references was confusing, are they there to state when the data was accessed?

Yes, the references are also now included in the References section correctly.

L272: missing comma: Before 2010, only CNEOS events Figure 3, label the y-axes as event count or number of events (consisent with f)) and number of phase picks for e)

Done.

L289: "About two-thirds of all detections occur between October and February" This was useful as I didn't notice any bias from the plot. A histogram with time of year on the x axis would be much clearer to verify the August sightings and winter frequencies. The length of day explaination would lead us to expect June/july to have the fewest sightings as these are the longest days around the solstice

We have adapted Figure 3d to show the month-event distribution in much more detail.

L292: "Perseid meteor shower" Do meteor showers have more fireballs? How was the IMO catalog filtered for fireballs, consistent with the the 0.3 kt threshold from the CNEOS catalog? Does the IMO observations include much smaller objects? describe the IMO fireball program and how they limit the reports to only include true fireballs and not ordinary meteors

We do not differentiate between meteors and fireballs. Our cutoff at 100 sightings per event is simply to limit the amount of data that was analyzed. In the future also events with less sightings should be investigated.

Line 294: missing comma "major showers, but instead are" or rearrange the syntax to not use a comma "major shower but are instead" How is the trajectory midpoint determined?

Done. We added an explanation to the midpoint calculation to line 151.

L305/306: "lateral distance of c. 160 km" I had to look up what c. meant

We keep ‘c.’, because it is widely used as abbreviation for circa.

L307 What offset? I don't understand what this sentence is saying. This whole paragraph is difficult to understand. I don't know how to explain the most likely distance is between 80 and 160 km. Pondering it on my own briefly, I think there could only be 1 station at 0 km but there is a large area of potential stations in the washer defined with inner ring at 80 km and outter at 160 km, beyond that, attenuation decreases detectability.

We added a more thorough explanation and also the expected decrease of phase picks due to decreasing signal strength over distances > 160 km.

L310: "acoustic emissions produced during the latter stages of flight at lower altitudes are expected to yield higher signal-to-noise ratios" It is not obvious where sound is most excited. There is regions in the trajectory where peak heating occurs and other regions where dynamic pressure is maximize. Of course, source-receiver distance and atmospheric ducts would also need considering making this a not straight forward understanding. A citation is necesary. I suggest Silber, E. A., & Bowman, D. C. (2025). Along-trajectory acoustic signal variations observed during the hypersonic re-entry of the OSIRIS-REx sample return capsule. Seismological Research Letters, 96(5), 2767-2779.

We added the reference as requested.

L334: Timing analysis seemed to come out of nowhere. I had to look back if I missed the introduction to it in the methods, but it only in the abstract.

We addressed this issue by describing the figure 4c in more detail.

L335: "Timing accuracy ... is likewise robust" I fail to see this demonstrated. How is the PN picker robust and how is the timing robust, just like it. robust against what changes, distance, impact energy, velocity, etc. If you are implying robust against the entirity of the dataset than that is rather subjective. Maybe 90% of results with 0.3 seconds is

This has been addressed by adding that the pick timing uncertainty is robust over the test subset and different meteoroid events.

L343: missing comma, compound sentence "reliably, but also" The results section seemed like a pile up. Restructure the Results section to SMORD database and interpretation of the distribution of data points and results of PN picker subsections Fig 4 c) pick uncertainty or timing error b) Define true positive rate, false positive rate. What is P, Maybe PN for PN picker model, maybe call the black line ROC "Histogram of the pick error in seconds versus the time accuracy" I think a clear concise phrasing describing this is a histogram of the timing error in seconds. Also it seems the words error, accuracy, and uncertainty are all used interchangedly concerning the timing of the onset. Also picks is used ambiguously. "Picks" is used to describe the manual identification of a signal onset, the PN picker labeling a P or N, and the time error. Establish a vocabulary for a reader to unambiguously understand. "Note the good match between both." The examples are a good match, but why pick the top 10 best examples. As none of these seem to be off by even 0.05 sec even though 90% of the data is within 0.3 sec and some are off by 10.0 sec as shown in panel c) I would like to see 10 random examples that would be more genuine.

We added an explanation for true positive rate, false positive rate, the ROC to the previous paragraph. We removed imprecise wording in the figure caption and switched from pick error to pick uncertainty. The samples we plotted in Figure 4d were chosen at random, but we include a bad example now as well.

L355: missing comma <meteoroid trajectory, which is> Fig 6. Define PSO

Done.

L381: "For applying infraGA we can incorporate" -> "By applying infraGA, we can incorporate"

Done.

L401: "For larger objects, the trajectory"

Done.

L403: "The data provided by SMORD ..." This sentence is repetitive. Two sentences prior, it was already stated SMORD could help with trajectories and ablation studies.

We have removed the repetition and moved the references into the previous paragraph.

L407: comparing amplitudes reveal information about the acoustic-to-seismic coupling efficiency, because amplitudes are useful in determining energy release. This doesn't logically follow. I don't know what this paragraph is for. Amplitude is useful for energy release, infrasound couples into the ground, there are more seismic stations than infrasound stations, so seismologic analysis should be advantageous for yield estimates. This sounds more like introductory content to justify using seismicity. The discusion section should have interpretations and implications of the results, uncertainty analysis, suggested improvements or next steps, etc.

We have moved this section to the evaluation of site effects and local coupling (lines 518ff) and reworded the paragraph to align more with a future application.

L415: "The search for acoustic-seismic signals can be quite difficult, because the recorded amplitude, and hence SNR, depends on numerous factors: source strength, distance or geometrical spreading, atmospheric winds, background noise, site effects, etc. Therefore, such signals are often not identified." Hence means "for this reason" and doesn't fit here, as writen. Are you trying to explain these signals are hard to find because low they often have low SNR?

We have modified the paragraph to be more concise and clearer about low-SNR meteoroid-related signals.

L418: "observers do often not recognize such waves" -> "observers often fail to recognize these waves"

Done. We also changed observers to seismic data analysts.

L420: a model is trained, not a dataset

This has been fixed.

L428: In our application, the PN predictions

Done.

L437: Here, we note

Done.

L492: there could be more data from IMO events with less than 100 observers (there could have been one observer awake at 3 am that witnessed an excellent example captured by many stations) and more data from CNEOS that include smaller yields

This is true. We included lowering the thresholds of both IMO and the CNEOS observations as a pathway to include more events into SMORD.