RC1: 'Comment on egusphere-2026-2100', Martin Murphy, 23 Jul 2026

1. line 163: “statistically robust”: I would interpret that to mean “can you randomly subset the towers and end up with the same general result?” Based on the spatial distribution in figure 6, it looks like the answer is “yes”. At the point of discussion of figure 6, it might be worth making an explicit comment in reference to those words from line 163

Answer: We thank the reviewer for this suggestion. We performed a bootstrap and random-subsampling analysis on the per-tower median LA values to test this directly. Resampling the 257 towers with replacement gives a 95% confidence interval on the median LA of [115.5, 135.2] m, consistent with the observed value of 124.9 m. Randomly subsampling 50%, 70%, and 90% of the towers (without replacement) confirms this stability: the 95% range of resulting medians narrows from ±10 m (50% subsample) to ±2 m (90% subsample) as sample size increases, indicating that the result is not driven by a small subset of towers.

We will add following reporting this analysis in the revised manuscript near the discussion of Fig. 6.: “To further support the robustness of the network-wide median LA, a bootstrap and random-subsampling analysis was performed on the 257 median LA values. Resampling with replacement (10,000 iterations) yields a 95% confidence interval of [115.5, 135.2] m around the observed median of 124.9 m. In addition, randomly subsampling 50%, 70%, and 90% of the towers (without replacement) produces median LA values that remain tightly clustered around the full-sample estimate, with the 95% range narrowing from ±10 m at 50% subsampling to ±2 m at 90% subsampling. This confirms that the reported median LA is not driven by a small subset of towers and supports the statistical robustness of the result.”

2. line 190: the question of an adaptive value of eps is a good one, but you could even further support the choice of 0.1 km by referring to figure 5a, where the ground-truth data from the 25-year history of Gaisberg tower validation and the DBSCAN analysis agree very nicely. That also adds more support to the question of "statistically robust".

Answer: Following will be added to the manuscript: Further support for this choice is provided by the cross-validation against the Gaisberg instrumented tower (Sect. 2.2.3, Fig. 5a), where the DBSCAN-based moving median LA, obtained using this fixed eps of 100 m, closely reproduces the ground-truth location accuracy over more than two decades, reinforcing both the appropriateness of this parameter choice and the statistical robustness of the methodology.

3. line 193/ figure 3: are the cluster sizes based on the entire period 2012-2025? I assume so, but it would probably be worth mentioning explicitly.

Answer: We confirm that the cluster sizes shown in Fig. 3 are derived from the DBSCAN clustering applied to the full 2012–2025 study period. We will make this explicit in the revised text “The spatial distribution of the towers and their associated cluster sizes, derived using the DBSCAN clustering algorithm applied to the full 2012–2025 study period, is shown in Figure 3.”

4. figure 7 / lines 276-294: I think that the "count" on the vertical axis of figure 7 is the number of structures, not the number of lightning events, but this probably should be made clear. Following from the discussion about the near-zero median and the definition of the 50% ellipse, the fact that the median offset is only -10 m says that almost exactly 50% of the towers have a larger median SMA than the DBSCAN results and 50% have smaller median SMA. That further supports the usefulness of the ellipse as a location accuracy metric where/when direct validation is not available. A potential "future work" analysis based on this paper, if one really wanted to dig deep into verifying the representativeness of the error ellipses, could be to look at the individual events that are included in

the DBSCAN clusters at a few of the towers, comparing their ellipses to their position offsets from the tower.

Answer: We thank the reviewer for this observation and for the additional interpretation, both of which we will incorporate into the revised manuscript.
Regarding the vertical axis of Fig. 7: the reviewer is correct that "Count" refers to the number of structures, not the number of individual lightning events. Each structure contributes exactly one value to this distribution, namely the signed difference between its DBSCAN-derived median LA and its median SMA (computed from all EUCLID events within 2 km of that structure). We will revise the figure caption to state this explicitly:

> "Figure 7: Distribution of signed differences between median LA (derived from DBSCAN clustering) and SMA values within 2 km of each structure. Each structure contributes one value to this distribution (N = 257); the count on the vertical axis therefore refers to the number of structures, not the number of individual lightning events.

We also thank the reviewer for pointing out the connection between the near-zero median offset and the definition of the 50% error ellipse. We propose adding the following to the discussion in Sect. 3.1:

> "Given that the median offset between the tower-based median LA and the SMA is close to zero (−10 m), this implies that for approximately half of the towers the median SMA exceeds the DBSCAN-derived LA, while for the other half it is smaller — consistent with the theoretical definition of the 50% error ellipse, in which the true strike location is expected to fall inside the ellipse in 50% of cases. This lends further support to the SMA as a practical proxy for location accuracy in the absence of direct ground-truth validation."

Finally, we agree that a valuable direction for future work would be to move beyond tower-level median comparisons and instead examine, for a selection of towers, the individual events within each DBSCAN cluster, comparing each event's own 50% error ellipse to its actual positional offset from the tower. We will add this suggestion to the Summary and Conclusions section as a proposed future extension of this methodology:

> “A further extension of this approach, left for future work, would involve examining individual events within the DBSCAN clusters of a selected subset of towers, comparing each event's own 50% error ellipse against its actual positional offset from the tower. Such an event-level analysis could help verify the representativeness of the error ellipse as a location accuracy metric in cases where direct ground-truth validation is unavailable.”