CC2: *One other comment of interest on the approach - a similar study on the NLDN was published in 2020 (https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2020JD032530) - how does the clustering approach they use compare with your DBSCAN approach?*

We thank you for this suggestion. We are aware of the Zhu et al. (2020) paper, since it is referenced at several places in the manuscript; however, we had not previously compared the two methods directly.

Zhu et al. (2020)'s method differs from ours in two aspects: (1) it applies a raster-based density threshold (median density over a 2 km domain, 50×50 m cells) with hard bounds on retained cluster size (25–250 cells), rather than a density-based point-clustering algorithm (DBSCAN) with a fixed neighborhood radius; and (2) when several density-elevated regions occur within the domain, Zhu's method keeps the one containing the domain-wide density maximum, whereas ours selects the cluster closest to the known structure location. The latter is only relevant when multiple high-density regions are detected, possibly due to a nearby structure not included in our dataset (e.g., an unlisted wind turbine).

Applying Zhu et al.'s original algorithm to the 257 structures for which our own method yielded a valid cluster, a qualifying cluster was found for 137 (53%; cf. 589/1451, 41%, for the NLDN in Zhu et al. 2020). For only 1 structure Zhu et al.’s method identified a cluster at ~1.5 km from the structure as a consequence of the domain-wide density maximum criterion. However, it is straightforward to modify Zhu's algorithm to instead retain the cluster closest to the tower, as in our method; doing so resolves this particular case (~90 m under both our own method and Zhu et. al.’s modified closest-to-tower rule) and gives a median difference (Zhu − ours) of −2.0 m (IQR 15.8 m) across the common structures for which a cluster was found under our closest-to-tower implementation, indicating the two methods are fundamentally consistent.

Note that Zhu's method does not find a qualifying cluster for all 257 towers. Further analysis shows this is largely due to Zhu's stricter cluster-size criterion (a retained cluster must span at least 25 grid cells): the fraction of towers for which Zhu's algorithm identifies a qualifying cluster at all rises from roughly 10% for towers where our own method finds a small cluster (10–50 events) to over 95% where our own cluster exceeds 250 events. We interpret this as a difference in minimum-data requirements between the two methods — Zhu's algorithm simply cannot return a result for the sparser clusters common in our dataset — rather than a disagreement about the underlying lightning locations. This is consistent with the wide range of lightning flash densities across our European dataset, from comparatively low densities in Scandinavia to substantially higher densities in Central and Southern Europe; our smaller minimum cluster size (min_samples = 10, vs. Zhu et al.'s 25-cell threshold) was chosen to remain applicable across this range.