EGU26-13365, updated on 13 Jun 2026
https://doi.org/10.5194/egusphere-egu26-13365
EGU General Assembly 2026

Joint Forecasting of Wind Speed and Direction via Ensemble Post-Processing

Alice Lake
United Kingdom of Great Britain – England, Scotland, Wales (alice.lake@metoffice.gov.uk)

As meteorological organisations transition to high-resolution ensemble-based forecasting, they risk leaving behind downstream users who rely on deterministic data: a need that may arise from the inability to process large volumes of data, or difficulty integrating probabilistic information into decision-making processes. Current solutions for such users typically involve providing the control (unperturbed) member of the ensemble, or deriving a single-value forecast through the independent treatment of variables (e.g., taking a median). However, relying solely on the control member discards the valuable information encoded within the full ensemble, fundamentally undermining the purpose of the ensemble. Meanwhile, univariate approaches can result in forecasts that lack physical consistency across variables. This limitation becomes critical when variables are interpreted jointly in real‑world decision‑making. Wind speed and direction exemplify this: these variables are used together in sectors such as renewable energy, where they inform turbine operation and resource planning, and aviation, where they underpin safety‑critical decisions around take‑off and landing. For these users, unrealistic combinations of speed and direction can translate directly into flawed risk assessments.

To address this gap, we present a novel ensemble post-processing technique that generates physically-consistent spot forecasts of wind speed and direction by exploiting the full ensemble distribution. The method constructs joint predictive probability density functions (PDFs) using a Gamma kernel for wind speed and a von Mises kernel for wind direction, accommodating the distinct statistical properties of these variables: non-negativity and skewness for speed, and circularity for direction. A single-value forecast is then obtained by selecting the ensemble member that maximizes its log-likelihood under the joint density across a specified forecast horizon. Because the selected forecast corresponds to one of the original ensemble members, it represents a physically plausible atmospheric state and maintains consistency across all variables, including those not directly analysed. This is critical for operational users: approaches that treat wind speed and direction separately (such as taking independent averages or applying separate post-processing to each variable) can produce unrealistic artefacts when passed through

downstream physical or statistical models.

This method was evaluated using the Met Office convective-scale ensemble, MOGREPS-UK, over the UK domain for a full calendar year, with verification at both the surface and aloft. Results are promising: the approach demonstrates the potential to outperform the control member, particularly at longer leadtimes where ensemble spread is greatest. These findings highlight an important step toward improving our offering to users and ensuring they remain supported as we transition to purely ensemble-based forecasting. Crucially, this work is not just theoretical; the next stage is to embed the technique into operational workflows and deliver it within user-facing products, ensuring these advances translate directly into improved real-world decision-making.