

Response to Editor:

Comment 1: Computational and Production Details: I recommend adding a concluding paragraph discussing the practical effort required to apply these models to a new site for short-term forecasting. Specifically, please detail the hardware requirements, training/evaluation times, and inference times. It would also be beneficial to briefly explain the codebase modifications necessary for production runs.

Reproducibility and README Updates: I reviewed the README file updated on December 9, 2025, and noted that it lacks step-by-step instructions for running the model. To improve reproducibility, please include a model run flowchart or pseudocode. Furthermore, please verify these instructions in a freshly installed environment and clearly document the coding environment details and specific module dependencies used.

Response: We sincerely thank Referee #1 for their final constructive comments. To ensure the highest standards of reproducibility and practical utility, we have made the following minor revisions prior to publication:

1. Computational and Production Details: We have added a concluding paragraph to the manuscript outlining the practical efforts required for deployment at a new site (Please refer to Page 22, lines 466-477 in the revised manuscript). This section details the specific hardware requirements, computational costs, and the necessary codebase modifications for production runs. For your convenience, the newly added text is quoted below:

“To apply the proposed framework to a new site in a production environment, the practical implementation is primarily divided into meteorological setup and operational inference. For the initial setup, running WRF requires a high-performance computing environment (e.g., 128-core CPU cluster), taking approximately 3 hours for a 72-hour forecast. Subsequent data processing, model training, and evaluation can be efficiently executed on standard consumer-grade hardware (e.g., an Intel Core i5-13600KF CPU, 16 GB RAM, and an NVIDIA RTX 2060 GPU). On such hardware, the GHI correction phase takes roughly 2 minutes, and the PINNs-iTransformer training takes approximately 3-5 minutes. Daily operational inference is computationally negligible (< 1 s). Applying this framework to a new region extends beyond routine codebase modifications, requiring

a rigorous, three-step re-determination process: (1) updating WRF domain coordinates and terrain files; (2) adapting the data ingestion pipeline to accommodate local PV specifications and weather observations; and (3) thoroughly retraining the QM-TPA-LSTM and PINNs-iTransformer architectures with local datasets. This mandatory retraining is crucial to ensuring that the framework’s physical consistency remains intact in any new environment.”

2. Reproducibility Enhancements: We comprehensively updated the project’s README file. It now includes step-by-step workflow instructions and explicit documentation of the tested software environment and dependencies. We also verified all instructions in a freshly installed environment to guarantee out-of-the-box reproducibility.

These updates fully address the referee’s requests and complete the development documentation for the MIPV-NWP-PINNs framework.

On behalf of all authors, best regards,

Xueshun Chen