Dear Editor,

We thank the editor and Anonymous Referee #2 for their careful examination of our source code and implementation. We have also submitted the revised version of our manuscript for your consideration.

To ensure transparency and reproducibility, we have uploaded the corrected source code and evaluation notebooks to our GitHub repository (`https://github.com/AMBehroozi/Flood_Neural_operator`). In addition, all datasets and pretrained model checkpoints have been deposited on Zenodo (`https://zenodo.org/uploads/23990987`). We have also updated the Code and Data Availability statement in the revised manuscript to provide access to these resources.

We strive to maintain a clean, well-organized, and reproducible research repository. However, during the iterative development process, temporary or outdated code may unintentionally remain in the repository. We continuously review and improve our publicly available code to identify and resolve such issues, remove unnecessary files, and ensure consistency with the final experimental workflow. We appreciate the reviewer's careful examination, which helps us further improve the clarity and reproducibility of our work.

Below, we provide our detailed clarifications and responses to the three source-code-related concerns raised by Anonymous Referee #2.

Comment 1: Residual Formulation of the FiLM-RCNO Magnifier

Response:

We thank the reviewer for carefully examining our implementation and providing this observation.

We respectfully refer the reviewer to **Figure 2 of the manuscript**, which illustrates our **nested, baseline-augmented residual formulation**. In this architecture, the bathtub baseline is combined with the learned convolutional output through an internal skip connection. The resulting Magnifier output is explicitly labeled *"Super-Resolution Depth Residual Correction"* in Figure 2, indicating that the hydraulic residual Δh represents the complete Magnifier output, rather than its final convolutional output alone.

The internal connection constructs this effective correction, while the external connection combines it with the bathtub baseline according to Equation (4).

Thus, although the baseline enters through two computational paths, the final reconstruction is expressed using **one baseline and one effective correction term**. This does not necessarily imply systematic overestimation of water depth, since the model is supervised using the complete high-resolution depth rather than a separate residual target, allowing it to compensate for both baseline contributions, as the reviewer correctly notes.

We confirm that this formulation is an integral part of our proposed methodology and was used to generate the reported results.

Comment 2: Initial Water-Depth Input During Evaluation

Response:

We thank the reviewer for identifying this inconsistency. We have corrected the evaluation code to use the actual initial water-depth fields, consistent with the training procedure, and rerun

all evaluation notebooks. **The correction resulted in only minor changes in the overall evaluation performance, without affecting the main findings or conclusions of the study.** These changes do not affect the main findings or conclusions of the study. The corresponding results in the manuscript have been updated accordingly.

Comment 3: Chowilla Inlet Hydrographs and Boundary Locations

Response:

We thank the reviewer for carefully examining our repository and bringing this issue to our attention. **Following this comment, we carefully reviewed the Chowilla dataset-generation workflow and confirmed that all datasets were generated using the correct gauge locations and configurations, with all three upstream inflow hydrographs applied at their intended boundary locations in the hydrodynamic simulations.** The inconsistency identified in the referenced notebook appears to have unintentionally remained from temporary modifications made during the development process. For example, the gauge locations mentioned by the reviewer appear to be remnants of the dam-break case study rather than the actual Chowilla simulation configuration. Such intermediate code may occasionally remain in a research repository even though it is not part of the final simulation workflow.

We appreciate the reviewer's careful observation, which helps us improve the clarity, consistency, and reproducibility of our research.

Sincerely,

Chaopeng Shen
Corresponding Author