XAI for Catchment Science – September 23-27, 2024

Instructor: Admin Husic, Associate Professor, Civil and Environmental Engineering, Virginia Tech
Contact: husic@vt.edu

Course Details

Meeting Time: Daily, 1-2:30 PM

- We will meet as a full group during this time.

Practical Session: Daily, 2:30-3:30 PM

- I will be available during this time for hands-on help and discussion for anyone who would like.

Meeting Location: Sala 2 PPGTA-UFMS

Course SharePoint: The course will be run through this SharePoint folder: [link TBD]

Computer Requirements: The course will involve a mixture of lectures, discussion, and breakout groups. Therefore, you will need to ensure you have a computer where you can look at and share materials.

Prerequisites:

- Some coding experience (Python, MATLAB, R, etc.). Code will be shared in the MATLAB and/or Python language. Generic datasets will be provided by the instructor.
- Participants are encouraged to prepare a dataset that they would like to explain with interpretable AI, preferably in a [x, y] array format. Here, "x" represents the number of sites/observations and "y" is the number of predictor variables. For each observation, there should be a target variable (the thing you want to predict, e.g., streamflow) and a matrix of predictor variables "y" (e.g., rainfall, forest cover, etc.).

Course Readings: Participants in the course should familiarize themselves with explainable AI techniques for decision tree models, such as random forest. In particular, we will leverage the Shapley value, and the following papers should be read ahead of time

1. Lundberg, Scott M., and Su-In Lee. "A unified approach to interpreting model predictions." *Advances in neural information processing systems* 30 (2017).
2. Lundberg, Scott M., et al. "From local explanations to global understanding with explainable AI for trees." *Nature machine intelligence* 2.1 (2020): 56-67.

Course overview

Course description

Linking how terrestrial drivers influence the quantity and quality of water in aquatic systems is a grand challenge of catchment science. This challenge grows more daunting as the effects of climate and land use change fundamentally shift catchment functions from their historic baselines to new norms. Models can be useful to explain system functions and make future predictions. To this end, black-box models provide unprecedented global predictive strength but lack clarity with respect to their local interpretability. Conversely, process-based models are locally interpretable but lack global transferability of insights.

This short-course will provide hydrology-oriented researchers with explainable artificial intelligence (XAI) tools that can be used to disambiguate local interactions from global drivers, thus simultaneously leveraging black-box predictive power and process-based interpretability. Given the recent surge in large-sample hydrology and catchment geochemistry, the advent of new machine learning techniques,

and the rapidly changing geophysical environment, this short-course will position researchers to enhance model interpretability and science advancement across many geoscience domains.

Students in the course will get hands-on experience with the development machine learning models such as those based on decision trees (e.g., random forest) for the prediction of continuous variables relevant to hydrology. XAI will then be applied to the machine learning models to extract global and local information from the model. Course participants will develop and explain models using their own data, provide peer-to-peer feedback, and leave the course with a toolset that can be used in future presentations and publications. A major goal of this course is to strengthen collaboration and knowledge exchange between local, regional, and global scientists in the effort to sustain our collective water resources.

Module learning objectives

1. Differentiate between process-based and data-driven modeling approaches.
2. Build random forest model to predict a target variable given a dataset of predictor features.
3. Explain machine learning models using interpretable AI techniques.
4. Discern between local explainability and global structure.
5. Create visualizations that can be used in future work including papers, reports, theses, etc.

Module schedule

Day	Topic	Activities and Exercises
September 23	Catchments as games	* Course overview * Datasets and scripts * Catchment predictors
September 24	Modeling fundamentals	* Building a random forest model * Evaluating model performance * Displaying model results
September 25	eXplainable AI (XAI)	* Explaining blackbox models * Partial dependence plots (global structure) * Shapley values (local structure)
September 26	Testing hypotheses	* Transform prediction to inference * Develop XAI model
September 27	Course presentations	* XAI model feedback

Assignments:

Each student will submit a "Course Report" by 3 PM on 10/01, which will be used for assessment. However, some of the items in this report will be used in class earlier. The course report will include:

- Message box assignment – SHARE IN CLASS 09/26
- XAI model creation assignment - 1 quantitative figure related to the students work

Assessment:

- XAI evaluation in-class presentations, 15%
- Message box in-class presentations, 15%
- In-class participation and enthusiasm, 20%
- Course report, 50%