


## Workshop: Accuracy with certainty: beyond the performance metrics

09/07/2024

Conferencia

•



- **Workshop (en inglés)**
- **Title:** Accuracy with certainty: beyond the performance metrics
- **Speaker:** María Moreno de Castro
- **Time and Date:** Martes 9 de julio, 11:30 - 13:30 hs.
- **Location:** Sala de Usos Múltiples, CITIC-UGR.
- **Organized by:** Centro de Investigación en Tecnologías de la Información y las Comunicaciones de la Universidad de Granada (CITIC-UGR).
- **Contact:** Rocio Celeste Romero Zaliz.

**Biography:** María Moreno de Castro, Theoretical Physicists by Autónoma de Madrid University, Master in Complex Systems Physics by Illes Balears University, and PhD in Kiel University with a thesis on Uncertainty Propagation. She currently works on the applications of Conformal Prediction for machine learning based time series analyses and co-organize PyData GRX meetups: <https://www.meetup.com/es-ES/pydatagr/>

**Abstract:** This workshop complements most supervised machine learning models applications, which often focus only on optimizing performance metrics. These metrics evaluate the quality of the model's predictions, but not how much we can trust them. For example, imagine a classifier trained to predict whether it rains (class 1) or not (class 0) given that day's temperature and humidity. Its performance is awesome, let us say its accuracy is >95% in validation and testing. Now, if the

<http://citic.ugr.es/>

prediction for a new day is '0.8', many books, courses, etc. conclude that 'there is an 80% probability of rain that day'. However, that raw score only went through a sigmoid that scaled it between 0 and 1 so all possible raw scores add up to 1, and it was not calibrated with the number of days that actually rained, therefore it is not a true probability. But it gets worse, what if we ask the model for a prediction where the input is the temperature and humidity of a day on the planet Venus? As the model has no way of telling us 'I have no idea what this is', it will just produce a biased prediction, without blushing. We can avoid both issues with Conformal Prediction (CP): it calibrates the predictions to be true probabilities, and, on the way, we obtain prediction intervals that, with mathematical guarantee, include the ground truth. The narrower the intervals, the greater the confidence in the prediction. CP is post-hoc (no model re-training is needed), model-agnostic (suitable for any model and task), distribution-free (there are zero parametric assumptions), and very lightweight (implemented with few lines of code, several packages are already developed, and run very fast). Several examples will be shown in the workshop, please find the materials in this repository:

[https://github.com/MMdeCastro/Uncertainty\\_Quantification\\_XAI](https://github.com/MMdeCastro/Uncertainty_Quantification_XAI).