

Title: Automated Machine Learning & Neural Architecture Search solutions, open challenges, and perspective

Instructor: Antonio Candelieri, University of Milano-Bicocca

Searching for accurate Machine/Deep Learning models, for a given dataset, is an awfully expensive process in terms of required computational resources and time. Automated Machine Learning (AutoML) is the process aimed at selecting, with a limited number of trials, the most performing ML/DL algorithm (i.e., Algorithm Selection) and optimally tuning its hyperparameters (i.e., Hyperparameter Optimization). With respect to Neural Networks, AutoML specializes in Neural Architecture Search (NAS), where architectural properties of the model are also optimized (e.g., number and types of layers). This course will provide the basics about Sequential Model Based Optimization – aka Bayesian Optimization – which is the de-facto standard methodology for AutoML and NAS.

Syllabus

- AutoML: Model Selection, Hyper-parameters Optimization and Combined Automated Model Selection and Hyper-parameters optimization (CASH)
- Search spaces properties and search strategies
 - Simple strategies: Grid Search vs Random Search
 - Sample efficient strategies: Sequential Model-based Bayesian Optimization
- Predicting the performance of ML algorithms: deterministic vs probabilistic models
- Dealing with prediction uncertainty: the exploration-exploitation trade-off
- AutoML for Neural Networks
 - Hyperparameter optimization and Neural Architecture Search (NAS)
- Hot research topics
- Software resources