



Challenge data



ENS

| PSL



COLLÈGE
DE FRANCE
—1530—

General documentation

This documentation contains explanations about:

- The challenges
 - Data
 - Score
 - Test set
 - Evaluation
- The leaderboards and their publications
 - Public leaderboard
 - Private leaderboard
- Schedule and key dates
 - New challenges
 - Award ceremony
- Getting started: creating an account

The challenges

Challenges are supervised learning problems (classification, regression, prediction, ranking, ...). From the input data that we write x , the goal is to predict a value, or an array of values y . This is done by an algorithm which associates to x a prediction $f(x)$. This prediction is learned from training data, where the input x and the output y are known.

The data

Every challenge contains at least 4 files:

- Train data: the data on which participants will elaborate and train their model, to learn the function f
 - o x_{train} : the inputs
 - o y_{train} : the outputs
- Test data: the data used for evaluation
 - o x_{test} : the inputs
 - o y_{test} : the outputs, this is the data participants shall predict, you don't have access to it

The score

Given an input x , the error between your prediction $f(x)$ and the true y is evaluated with a score function l which returns a real number $l(f(x), y)$. Your goal is to minimize the average score on the inputs x_i . The calculation of the score depends upon the challenge and is explained in the presentation of each challenge. It may be based on a Mean Square Error, a Cross Entropy, an R2 score, AUROC A custom loss can be specified in some challenges.

Test set separation and overfitting

The test set is separated into two halves for the evaluation, the public and the private test sets. Every submission by a participant gives two scores, one on each subset.

- public score: the score computed on the public set. It is provided at each submission.
- private score: the score obtained on the private set. It is provided twice a year.

This separation prevents participants from overfitting the whole test set, as they can't see how their solution behaves on the entire test set. To limit overfitting on the public set, the number of submissions is **limited to 2 per day** for each participant. Every time a participant submits a solution, he gets his score on the public set. Participants must train an algorithm that generalizes on both the public set and the private set. The private score is given every year, on the 15th of June and the 15th of December.

The leaderboards and their publications

Participants can choose which submission is used for the leaderboards. By default the submission with highest score on the public set is chosen. From submissions selected by each participant we establish the following leaderboards.

Public leaderboard

It ranks the submission scores of participants on the public test set. It is **permanently** available.

Private leaderboard

It ranks the submission scores on the private test set. It is **hidden** to participants. The private leaderboard is only revealed twice a year: on **June 15th** and **December 15th**.

Schedule and key dates

Challenges are continuously running but special events are organized every year, to which anyone is welcome to attend.

New challenges

In January of every year, a new season of challenges is added to the platform, on top of the previous seasons. Sessions of presentations of these new challenges are organized at the Collège de France, as part of the course of Professor Stéphane Mallat. See for example the presentations of challenges in 2020 <https://www.college-de-france.fr/site/stephane-mallat/Challenges-2020.htm>.

Award ceremony

We celebrate the winners of the challenges of the previous year, at the beginning of February. These winners have obtained the best score on the December 15th **private leaderboard**. The prizes are thus awarded at the end of the first year of competition.

Getting started: creating an account

To create an account, click on the “register” tab. There are three types of accounts:

- participant
- professor
- provider

When creating your account, make sure that your email address is not already used. You can't have two accounts for the same email.

For professors or providers, your account must be validated by us. For professors, make sure to fill the “institution” field. Your validation will be easier if you use a professional address from your institution.

Supported by

