

Data Mining Project: *Tennis Matches*

A **project** consists in data analysis based on the use of data mining tools. The project has to be performed by a team of 2/3 students. It has to be performed by using Python. The guidelines require to address specific tasks and results must be reported in a unique paper. The total length of this paper must be **max 25 pages** of text including figures. The students must deliver both: paper and well commented Python notebooks.

DATASET DESCRIPTION

The dataset to be analyzed is composed of information about some tennis matches. Data are organized in .csv file:

- tennis_matches.csv where each row represents a match
- male_players.csv
- female_players.csv

The dataset contains different features, with the following semantic meaning:

- **Tourney_id**: a unique identifier for each tournament, such as 2020-888. The exact formats are borrowed from several different sources, so while the first four characters are always the year, the rest of the ID doesn't follow a predictable structure.
- **Tourney_name**: the name of the tourney
- **Surface**: kind of surface for the match
- **Draw_size**: number of players in the draw, often rounded up to the nearest power of 2. (For instance, a tournament with 28 players may be shown as 32.)
- **Tourney_level**: they are split for men and women.
 - **For men**: 'G' = Grand Slams, 'M' = Masters 1000s, 'A' = other tour-level events, 'C' = Challengers, 'S' = Satellites/ITFs, 'F' = Tour finals and other season-ending events, and 'D' = Davis Cup. F
 - **For women**, there are several additional tourney_level codes, including 'P' = Premier, 'PM' = Premier Mandatory, and 'I' = International. The various levels of ITFs are given by the prize money (in thousands), such as '15' = ITF \$15,000. Other codes, such as 'T1' for Tier I (and so on) are used for older WTA tournament designations. 'D' is used for the Federation/Fed/Billie Jean King Cup, and also for the Wightman Cup and Bonne Bell Cup.
 - There is also some competition which can be for both men and women: 'E' = exhibition (events not sanctioned by the tour, though the definitions can be ambiguous), 'J' = juniors, and 'T' = team tennis, which does yet appear anywhere in the dataset but will at some point.
- **Tourney_date**: eight digits, YYYYMMDD, usually the Monday of the tournament week.
- **Match_num**: a match-specific identifier. Often starting from 1, sometimes counting down from 300, and sometimes arbitrary.
- **Winner_id**: the player_id used in this repo for the winner of the match.

- **Winner_entry:** 'WC' = wild card, 'Q' = qualifier, 'LL' = lucky loser, 'PR' = protected ranking, 'ITF' = ITF entry, and there are a few others that are occasionally used.
- **Winner_hand:** R = right, L = left, U = unknown. For ambidextrous players, this is their serving hand.
- **Winner_ht:** height in centimetres, where available
- **Winner_ioc:** three-character country code
- **Winner_age:** the age of the player, in years, depending on the date of the tournament
- **Best_of:** '3' or '5', indicating the number of set for this match
- **Minutes:** match length, where available
- **W_ace:** winner's number of aces
- **W_df:** winner's number of doubles faults
- **W_svpt:** winner's number of serve points
- **W_1stIn:** winner's number of first serves made
- **W_1stWon:** winner's number of first-serve points won
- **W_2stwon:** winner's number of second-serve points won
- **W_SvGms:** winner's number of serve games
- **W_bdSaved:** winner's number of breakpoints saved
- **W_bdFaced:** winner's number of breakpoints faced
- **Winner_rank:** winner's ATP or WTA rank, as of the `tourney_date`, or the most recent ranking date before the `tourney_date`
- **Winner_rank_points:** number of ranking points, where available.

We did not report the meaning of losers attribute because they are the same as the winners, but the feature names start with 'loser'.

Task 1 Data Understanding and Preparation (30 points):

Task 1.1: Data Understanding: Explore the dataset with the analytical tools studied and write a concise "data understanding" report assessing data quality, the distribution of the variables and the pairwise correlations.

Task 1.2: Data Preparation: Improve the quality of your data and prepare it by extracting *new features* interesting for describing the player profile and his behavior derivable from matches. These indicators have to be extracted for each player.

Examples of Indicators to be computed are:

- how many times did the player win during a given period
- how many matches the player played in a given period
- a ratio between the previous indicators
- percentage of aces related to the number of first serves made
- number of breakpoints numbers w.r.t. all games
-

Note that these examples are not mandatory. You can derive indicators that you prefer and that you consider interesting for describing the players.

It is MANDATORY that each team defines **indicators** and their description and when it is necessary also their mathematical formulation. The profile will be useful for the clustering analysis (i.e., the second project's task).

Once the set of indicators is computed, the team has to explore the new features for a statistical analysis (distributions, outliers, visualizations, correlations).

Subtasks of DU

- Data semantics for each feature that is not described above and the new one defined by the team
- Distribution of the variables and statistics
- Assessing data quality (missing values, outliers, duplicated records, errors)
- Variables transformations
- Pairwise correlations and eventual elimination of redundant variables

Task 2: Clustering analysis (30 POINTS - 32 with optional subtask)

Based on the player's profiles explore the dataset using various clustering techniques. Carefully describe your decisions for each algorithm and which are the advantages provided by the different approaches.

Subtasks

- Clustering Analysis by K-means:
 1. Identification of the best value of k
 2. Characterization of the obtained clusters by using both analysis of the k centroids and comparison of the distribution of variables within the clusters and that in the whole dataset
 3. Evaluation of the clustering results
- Analysis by density-based clustering:
 1. Study of the clustering parameters
 2. Characterization and interpretation of the obtained clusters
- Analysis by hierarchical clustering
 1. Compare different clustering results got by using different merging strategies
 2. Show and discuss different dendrograms using the different merging strategies
- Final evaluation of the best clustering approach and comparison of the clustering obtained
- **Optional (2 points):** Explore the opportunity to use alternative clustering techniques in the library: <https://github.com/annoviko/pyclustering/>

Delivery of the first draft of the report with Task 1.1, Task 1.2 and Task 2: 5 November

Task 3: Predictive Analysis (30 POINTS)

Consider the problem of predicting for each player a label that defines if s(he) is a *high ranked player* or a *low ranked player* (binary task) by exploiting the feature related to the rank of the players.

The student need to:

1. Define a player profile that enables the above player classification. For this task, you can exploit the profile created for the clustering task, by adding or removing features, depending on the results previously obtained.
2. Compute the label for any customer. The extraction of the label can take advantage of several features related to the rank, such as `loser_rank`, `winner_rank`, `loser_rank_points`, `winner_rank_points`, etc. An example of simple label can be derived by:
 - computing the average rank per player by exploiting `loser_rank`, `winner_rank`
 - select a threshold for discretizing in two categorical labels the class.

Note that you can define in different ways the labels.

3. Perform the predictive analysis comparing the performance of different models, discussing the results and discussing the possible preprocessing that you applied to the data for managing possible problems identified that can make the prediction hard. Note that the evaluation should be performed on both training and test set.

Note: The final report delivered within the end of December can also improve the already delivered tasks.

Task 4: Address one of the two tasks (32 POINTS - Optional)

Task 4.1: Time series analysis

Consider the dataset of time series `CityGlobalTemperature2000-2009.csv` containing for 100 cities the temperature measurements (mean and standard deviation over a month). The goal of the task is to find groups of similar cities with respect to the temperature trends.

Task 4.2: Explanation Analysis

Consider the non-interpretable models used in Task 3 (eg. SVM, ensemble methods, etc) and study the global explanation with SHAP and the local explanation with LIME and SHAP. Use the evaluation metrics presented during the XAI Laboratory (for using LORE, you can download the library at the link: [https://github.com/rinziv/XAI lib HAI-net Tutorial](https://github.com/rinziv/XAI_lib_HAI-net_Tutorial) and follow the instructions contained in the notebook presented during the XAI Laboratory).

Rules for final delivery and Exam

Project Delivery. The final deadline of the project is **5th January 2021 at 23:59**. This deadline is **STRICT**. No extension is possible because then the winter session of exams starts.

Each group must deliver by email to anna.monreale@unipi.it and francesca.naretto@sns.it a zipped folder named **DM_GroupID.zip** and containing 4 folders and 1 pdf file:

1. a folder named **DM_GroupID_TASK1**, containing source code of data understanding
2. a folder named **DM_GroupID_TASK2**, containing source code of data clustering
3. a folder named **DM_GroupID_TASK3**, containing source code of classification
4. (optional) a folder named **DM_GroupID_TASK4**, containing source code of time series analysis/explanation analysis
5. a pdf file with maximum 25 pages including figures discussing the results of the 4 tasks. The name of this file must be: **DM_Report_GroupID.pdf**. The file must contain the list of authors (i.e., members of the group).

Remember that the final submission can contain updated versions of the work already delivered in the previous deadlines.

ATTENTION: On 5th January 2022, I will publish the text of the new project description for students who will not deliver the project within the fixed deadline.

Exam

There are two possible options for the exam:

1. project presentation + questions on the whole program
2. project presentation + paper presentations (in the dates already fixed)

I prefer to have group presentations of the project. If this is impossible otherwise we can find a solution together.

How to book for the exam colloquium?

Soon in <https://esami.unipi.it/> you will find the dates for the exam (they should be 17 Jan and 9 Feb).

Each student must do the registration on one of the 2 dates. These are not the dates of the colloquium but I will use the list of students for organizing the exam dates.

Final Grade

The final grade of the exam is given by the weighted average of the project evaluation and oral/paper presentation evaluation. I will assign a weight of 70% to the project work and 30% to the oral/paper presentation. Consider that the project evaluation also includes the project presentation (my suggestion is using slides). The contribution of each task to the project evaluation will be: 35% for the Task1, 30% Task2, 30% Task3 and only 5% Task4 that is optional.

Remember that any student must be able to answer any question on the project work.

I will not communicate any numerical evaluation for the project before the oral exam. I will only communicate if the evaluation is positive or negative 3 days before 17 Jan.