Intelligent Disease Progression Prediction: Overview of iDPP@CLEF 2022

Alessandro Guazzo[1,*], Isotta Trescato[1,*], Enrico Longato[1], Enidia Hazizaj[1], Dennis Dosso[1], Guglielmo Faggioli[1], Giorgio Maria Di Nunzio[1], Gianmaria Silvello[1], Martina Vettoretti[1], Erica Tavazzi[1], Chiara Roversi[1], Piero Fariselli[2], Sara C. Madeira[3], Mamede de Carvalho[3], Marta Gromicho[3], Adriano Chiò[2], Umberto Manera[2], Arianna Dagliati[4], Giovanni Birolo[2], Helena Aidos[3], Barbara Di Camillo[1], and Nicola Ferro[1]

[1] University of Padua, Italy
{barbara.dicamillo, giorgiomaria.dinunzio, dennis.dosso, nicola.ferro, enrico.longato, gianmaria.silvello, martina.vettoretti, erica.tavazzi}@unipd.it, {guglielmo.faggioli, alessandro.guazzo, isotta.trescato}@phd.unipd.it, {enidia.hazizaj, chiara.roversi}@studenti.unipd.it
[2] University of Turin, Italy
{adriano.chio, giovanni.birolo, piero.fariselli, umberto.manera}@unito.it
[3] University of Lisbon, Portugal
mamedemg@mail.telepac.pt, mgromichosilva@medicina.ulisboa.pt, {haidos, sacmadeira}@fc.ul.pt
[4] University of Pavia, Italy
arianna.dagliati@unipv.it

Abstract. *Amyotrophic Lateral Sclerosis (ALS)* is a severe chronic disease characterized by progressive or alternate impairment of neurological functions, characterized by high heterogeneity both in symptoms and disease progression. As a consequence its clinical course is highly uncertain, challenging both patients and clinicians. Indeed, patients have to manage alternated periods in hospital with care at home, experiencing a constant uncertainty regarding the timing of the disease acute phases and facing a considerable psychological and economic burden that also involves their caregivers. Clinicians, on the other hand, need tools able to support them in all the phases of the patient treatment, suggest personalized therapeutic decisions, indicate urgently needed interventions. The goal of iDPP@CLEF is to design and develop an evaluation infrastructure for AI algorithms able to:

1. better describe disease mechanisms;
2. stratify patients according to their phenotype assessed all over the disease evolution;
3. predict disease progression in a probabilistic, time dependent fashion.

[*] These authors contributed equally.

1 Introduction

Amyotrophic Lateral Sclerosis (ALS) is a neurological disease that causes the progressive degeneration of the motor neurons that control voluntary muscles, resulting in an increasing impairment of motor and vital functions and leading to death usually within 4-5 years from the diagnosis. Likely resulting from a complex interplay of genetic and environmental factors, ALS is characterized by high heterogeneity in both symptoms and disease progression, especially in the early stages of the disease. This heterogeneity is partly responsible for the lack of effective prognostic tools in medical practice, as well as for the current absence of a therapy able to effectively slow down or reverse the disease course. On the one hand, patients need support for facing the psychological and economic burdens deriving from the uncertainty of how the disease will progress; on the other, clinicians require tools that may assist them throughout the patient's care, recommending tailored therapeutic decisions and providing alerts for urgently needed actions.
In order to improve the current diagnostic and prognostic situation, we should design and develop *Artificial Intelligence (AI)* algorithms be able to:

- stratify patients according to their phenotype, assessed all over the disease evolution;
- predict the progression of the disease in a probabilistic, time dependent fashion;
- better describe disease mechanisms.

The *Intelligent Disease Progression Prediction at CLEF (iDPP:CLEF)* lab[5] aims to design and develop an evaluation infrastructure for driving the development of such AI algorithms. By "evaluation infrastructure", we mean experimental collections, evaluation protocols, evaluation measures, ground-truth creation protocols, and so on. Indeed, in this context, it is fundamental, even if not so common yet, to develop shared approaches, promote the use of common benchmarks, and foster the comparability and replicability of the experiments. Differently from previous challenges in the field, iDPP:CLEF addresses in a systematic way some issues related to the application of AI in clinical practice in ALS. Therefore, in addition to defining the risk scores based on the probability that an event will occur in the short or long term period, iDPP:CLEF also addresses the issue of providing information in a more structured and understandable way to clinicians.

The paper is organized as follows: Section 2 presents related challenges; Section 3 describes its tasks; Section 4 discusses the developed dataset; Section 5 explains the setup of the lab and introduces the participants; Section 6 introduces the evaluation measures adopted to score the runs; Section 7 analyzes the experimental results for the different tasks; finally, Section 8 draws some conclusions and outlooks some future work.

[5] https://brainteaser.health/open-evaluation-challenges/idpp-2022/

2 Related Challenges

To the best of our knowledge, within CLEF, there have been no other labs on this or similar topics before.

Outside CLEF, there have been a recent challenge on Kaggle[6] in 2021 and some older ones, the DREAM 7 ALS Prediction challenge[7] in 2012 and the DREAM ALS Stratification challenge[8] in 2015.

The Kaggle challenge used a mix of clinical and genomic data to seek insights about the mechanisms of ALS and difference between people with ALS who progress faster versus those who develop it more slowly. The DREAM 7 ALS Prediction challenge [12] asked to use 3 months of ALS clinical trial information (months 0–3) to predict the future progression of the disease (months 3–12), expressed as the slope of change in *ALS Functional Rating Scale Revisited (ALSFRS-R)* [5], a functional scale that ranges between 0 and 40. The DREAM ALS Stratification challenge asked participants to stratify ALS patients into meaningful subgroups, to enable better understanding of patient profiles and application of personalized ALS treatments.

Differently from these previous challenges, iDPP:CLEF focuses on explainable AI and on temporal progression of the disease.

3 Tasks

iDPP:CLEF 2022 is the first edition of the lab and consists of pilot activities aimed both at an initial exploration of ALS progression prediction and at understanding of the challenges and limitations to refine and tune the labs itself for future iterations.

In particular, iDPP:CLEF targets two kinds of activities:

1. preliminary and exploratory pilot tasks on disease progression prediction;
2. position papers on the explainability of the prediction algorithms.

Overall, this mix provides participants with the opportunity to make some hands-on experience with these data and provide feedback about the task design as well as to brainstorm on how to evaluate this kind of algorithms and, in particular, assess their explainability.

3.1 Pilot Task 1: Ranking Risk of Impairment

As shown in Figure 1, this task focuses on ranking of patients based on the risk of impairment in specific domains. More in detail, we use the ALSFRS-R scale to monitor speech, swallowing, handwriting, dressing/hygiene, walking and respiratory ability in time and ask participants to *rank patients based on time to event risk* of experiencing impairment in each specific domain.

[6] https://www.kaggle.com/alsgroup/end-als
[7] https://dreamchallenges.org/dream-7-phil-bowen-als-prediction-prize4life/
[8] https://dx.doi.org/10.7303/syn2873386.

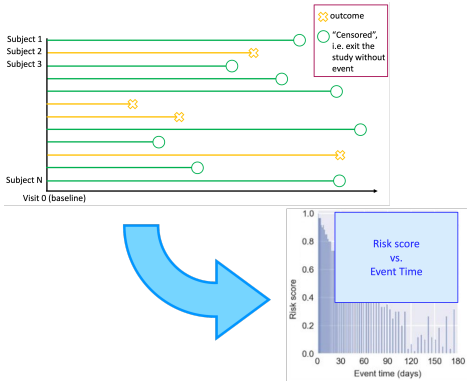

Fig. 1: Task 1: from patients to ranking of patients based on time of event risk.

More in detail, participants are asked to rank subjects based on the risk of early occurrence of

- **Task 1a**: *Non-Invasive Ventilation (NIV)* or (competing event) Death[9], whichever occurs first;
- **Task 1b**: *Percutaneous Endoscopic Gastrostomy (PEG)* or (competing event) Death, whichever occurs first;
- **Task 1c**: Death[10].

For each of these tasks, participants are given a dataset containing 6 months of visits and are asked to rank patients on the risk of occurrence of one of the above events after month 6.

In particular, for each sub-task, we ask for two types of submission from participants:

- submissions using only data available until `Time 0`, i.e. the time of the first ALSFRS-R questionnaire;
- submissions using data available until `Month 6`.

Indeed, from the clinicians point of view, it is of interest to understand what they can say the first time they see the patient (`Time 0`) and what they can say if they collect additional data for the following 6 months.

3.2 Pilot Task 2: Predicting Time of Impairment

As shown in Figure 2, this task refines Task 1 asking participants to *predict when specific impairments will occur* (i.e. in the correct time-window). In this regard,

[9] Death is considered a competing event since a patient might incur death before experiencing the event of interest; the models should account for that.

[10] For the tasks 1c and 2c, death is not a competing event anymore but the focus of the models' predictions.

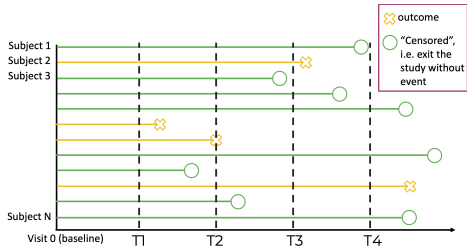

Fig. 2: Task 2: from patients to time of impairment.

we assess model calibration in terms of the ability of the proposed algorithms to estimate a probability of an event close to the true probability within a specified time-window.

In particular, participants are asked to predict the time of the event. Where the event is

- **Task 2a**: NIV or (competing event) Death, whichever occurs first;
- **Task 2b**: PEG or (competing event) Death, whichever occurs first;
- **Task 2c**: Death.

As in the previous case, for each sub-task, we ask two type of submissions from participants:

- submissions using only data available until `Time 0`, i.e. the time of the first ALSFRS-R questionnaire;
- submissions using data available until `Month 6`.

3.3 Position Papers Task 3: Explainability of AI algorithms

This task is not an evaluation challenge but rather a discussion on how to make these prediction algorithms explainable, also in a visual way.

Therefore, this task called for position papers to start a discussion on AI explainability including proposals on how the single patient data can be visualized in a multivariate fashion contextualizing its dynamic nature and the model predictions together with information on the predictive variables that most influence the prediction. We evaluated proposals of different visualization frameworks able to show the multivariate nature of the data and the model predictions in an explainable, possibly interactive, way.

Even if this task is not an evaluation challenge, authors of the papers were welcome to use the datasets provided by iDPP:CLEF, if they wished to give examples of their algorithms and solutions, or to explore the submissions made by other participants in iDPP:CLEF and apply their explainability techniques to them.

Table 1: Main features of the iDPP:CLEF dataset.

Section	Sub-section	Variables
Baseline	Patient	Sex, Date of Birth
	ALS Onset	Date, Site
	Diagnosis	Date, Regions affected, Diagnostic Delay, FVC, BMI at diagnosis
Follow-up	Progression scores	ALSFRS-R, Rate of disease progression
	Tests	Hematologic tests, Muscle strength assessed by manual testing, Respiratory function tests
	Therapy	ALS treatments
	Other	Regions affected, Upper and lower motor neuron signs, Cognitive and neurophysiological changes
Clinical Events	History	BMI premorbid, Family history, Comorbidities, Previous surgery and trauma
	Interventions	Date of NIV, Date of PEG, Date of Tracheostomy
	Survival	Date of death
Lifestyle	Lifestyle	Working activity, Physical activity, History of smoking, Marital status, Education level

4 Dataset

iDPP:CLEF developed a dataset containing patient records from two clinical institutions in Turin, Italy, and in Lisbon, Portugal.

The dataset is fully anonymized, meaning that all the information which might reveal the identity of a patient, e.g. place of birth or city of residence, are removed; we also avoided absolute dates and made everything relative to `Time 0`, i.e. the date of the first ALSFRS-R questionnaire [5].

Table 1 summarizes the main features and variables available in the dataset. The following data are available for both the training and the test sets:

- the first available ALSFRS-R questionnaire at `Time 0` (both single question scores and total score).
  Thus, for example, time-of-onset and time-of-diagnosis are expressed as relative delta with respect to `Time 0` in months (also fractions);
- the slope of the ALSFRS-R score between time-of-onset and `Time 0` as:

$$slope = \frac{48 - \text{ALSFRS-R-score}\,(\texttt{Time 0})}{\texttt{Time 0} - \texttt{TimeOnset}}$$

- all the other static data, with a complete list available at http://brainteaser.dei.unipd.it/challenges/idpp2022/assets/other/static-vars.txt
- visits , containing either other ALSFRS-R questionnaires or Spirometry, i.e. *Forced Vital Capacity (FVC)*. The complete list of variables for each visit

is available at http://brainteaser.dei.unipd.it/challenges/idpp2022/assets/other/visits.txt.

We ensured that, for each patient, there are 6 months of data, so that predictions can be made using either only data available at `Time 0` or all the data available until month 6.

The following data are available only for the training set:

- Time of event (`NIV`, `PEG`, or `DEATH`); or
- Censoring time, i.e. time of the last available visit if none of the previous events occurs;

according to the following format:

```
0x4bed50627d141453da7499a7f6ae84ab 1 PEG 20.5
0x4d0e8370abe97d0fdedbded6787ebcfc 1 PEG 18.3
0x5bbf2927feefd8617b58b5005f75fc0d 1 DEATH 17.6
0x814ec836b32264453c04bb989f7825d4 0 NONE 37.4
0x71dabb094f55fab5fc719e348dffc85 1 PEG 8.2
...
```

where:

- Columns are separated by a white space;
- The first column is the `patient ID`, a 128 bit hex number (should be considered just as a string);
- The second column indicates whether the one of the above events occurred (`1`) or not (`0`);
- The third column is the occurred event. It comes from a controlled vocabulary and it can be either `NIV`, `PEG`, `DEATH`, or `NONE`;
- The fourth column is the time of the event, or the censoring time, from `Time 0` in months.

Training and test datasets follow a (roughly) 80%-20% proportion; more details about the split into training and test are provided below.

Both Task 1 and Task 2 use the same datasets but we prepared a separate dataset for each of the sub-tasks to make it simpler for participants to focus on a specific event to be predicted. Table 2 provides details about the created datasets.

Creation of the datasets The full dataset contained approximately 4,800 records linked to patients, with around 20,000 ALSFRS-R questionnaires in total and 5,500 records concerning spirometries. The original data contain minor inconsistencies and typos. Therefore, we first process the data, removing records that are likely wrong or do not provide essential information to enable prediction. In terms of patient records we removed those presenting an unordered sequence of events (i.e., onset after diagnosis or diagnosis after death). Such event sequences are likely due to typos and other human errors, which result in wrong records that might introduce noise and spurious information in the final dataset.

Table 2: Training and test datasets.

Training				
Sub-task	**Patients**	**ALSFRS-R**	**Spirometry**	**Outcome**
Sub-task a	1,454	3,668	1,189	– `NIV`: 675 patients (46.42%) – `DEATH`: 636 patients (43.74%) – `NONE`: 143 patients (9.83%)
Sub-task b	1,715	4,264	1,506	– `PEG`: 501 patients (29.21%) – `DEATH`: 969 patients (56.50%) – `NONE`: 245 patients (14.29%)
Sub-task c	1,756	4,366	1,536	– `DEATH`: 1,486 patients (84.62%) – `NONE`: 270 patients (15.38%)
Test				
Sub-task	**Patients**	**ALSFRS-R**	**Spirometry**	**Outcome**
Sub-task a	350	872	273	– `NIV`: 162 patients (46.29%) – `DEATH`: 152 patients (43.43%) – `NONE`: 36 patients (10.29%)
Sub-task b	430	1,049	361	– `PEG`: 120 patients (27.91%) – `DEATH`: 251 patients (58.37%) – `NONE`: 59 patients (13.72%)
Sub-task c	494	1,220	414	– `DEATH`: 417 patients (84.41%) – `NONE`: 77 patients (15.59%)

Furthermore, a patient record was dropped if one or more of the following pieces of information were absent:

– onset or diagnosis dates;
– death date in records associated with dead patients;
– at least six months of historical ALSFRS-R questionnaires before an event (NIV, PEG, or (competing event) Death).

We adopt the filtering strategy mentioned above to grant that every record in the final dataset contains enough information to allow proper predictions.

Concerning the ALSFRS-R questionnaires, we removed those records that had one or more of the following problems:

– duplicate records;
– missing date;
– one or more of the ALSFRS-R items missing;

– *ALS Functional Rating Scale (ALSFRS)* questionnaires with the old formulation (thus with items from 1 to 9, plus the old 10^{th} item). We include only records referring to ALSFRS-R.

Furthermore, if one or more of the ALSFRS-R sub-scores or the total ALSFRS-R score do not agree with the sum of the associated ALSFRS-R items, we replace the value reported in the original dataset with the sum of the linked items. Finally, regarding the spirometries, we removed duplicated records, records with a missing date, and FVC percentage value.

Figure 3 illustrates a set of - synthetic - patients and their clinical history, describing whether they satisfy the conditions to be inserted into the dataset. By construction, the first ALSFRS visit (blue bullets) is considered as `Time 0`, while the moment of the previous spirometries (yellow bullets) and subsequent visits is indicated as the difference in months with respect to the reference ALSFRS.

– Patient 1 is inserted into the dataset, having a proper sequence of visits, questionnaires and events (at least six months of information before the first event).
– Patient 2, on the other hand, cannot be included in the dataset since they do not have enough information.
– For Patient 3, we observe that only four months passed between the first ALSFRS and the first event. Thus, even though we have 6 months of overall information (first spirometry to event), we cannot retain the record.
– Patient 4, regardless of the fact that they have a single ALSFRS, can be included in the dataset since the distance between the first ALSFRS and the event is above six months.
– Both patients 5 and 6 need to be excluded from further analyses: the former does not have six months of information before the first event, while the latter does not have enough history, regardless of the spirometry taken before the first ALSFRS.
– Patients 7 and 8, on the other hand, can be considered: the former has a proper clinical history, while the latter, even though he or she have a "censoring" event – marked with a question mark, has more than six months of history.

Split into training and test Each of the three available datasets (sub-task a, b, and c) was split into a training set and a test set, with proportions 80% and 20%, respectively. The data were split stratifying the subjects according to outcome time and to the specific outcome type (*death, NIV, none* for sub-task a, *death, PEG, none* for sub-task b, and *death, none* for sub-task c). Stratifying by these two variables is instrumental to the fairness of the challenge as it forces an equal distribution of their levels across the two subsets. The simplest method to verify whether stratification has been performed correctly is to compare the distribution of the stratification variables (*outcome time* and *outcome type*) in each training/test pair. From the literature, certain variables are known to be

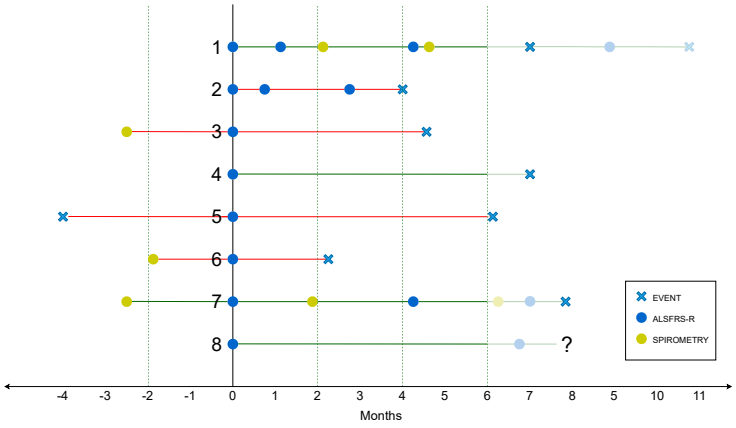

Fig. 3: Sequences of events that allow (or forbid) a patient to be considered as suitable to belong to the dataset. Events grayed out refer to those events happened after another and thus ignored. Visits grayed out refer to visits happened after the first six months.

particularly relevant in predicting events related to ALS progression [6], therefore, even though they were not included in the stratification criteria, we verified that *sex*, *age at onset*, *onset site*, *ALSFRS-F slope*, and the number of available visits in the first 6 months were also equally represented in the training and test sets. Table 3 reports, as an example, the comparison of the variables' distributions in the training (second column) and test (third column) sets for sub-task a. The comparison for the other two sub-tasks are shown in the extended overview [8]. Since the distributions are similar, we concluded that the training/test split provided to the participants met best-practice quality standards.

5 Lab Setup and Participation

In the remainder of this section, we detail the guidelines the participants had to comply with to submit their runs and the submissions received by iDPP@CLEF.

5.1 Guidelines

Participating teams were provided with the following guidelines:

- The runs should be submitted in a textual format in the participant repository, both described below;
- Each group can submit a maximum of 5 runs for each sub-task, thus amounting to maximum 15 runs for each of Task 1 and Task 2;
- For each task, participants are asked to submit two types of runs: either using only the information available at `Time 0` or using all the information available in the first 6 months.

Table 3: Sub-task a, comparison between training and test populations. Continuous variables are presented as *median [1st - 3rd quartiles]*; discrete variables as *count (percentage on sample total)*, for each level.

	Training	**Test**
Number of subjects	1454	350
Outcome type	Death: 636 (44%) NIV: 675 (46%) Censoring: 143 (10%)	Death: 152 (43%) NIV: 162 (46%) Censoring: 36 (10%)
Outcome time	17.75 [11.14-30.99]	20.72 [11.25-36.76]
Sex	M: 743 (51%) F: 711 (49%)	M: 188 (54%) F: 16 (46%)
Age at onset	64.89 [55.66-70.76]	64.76 [56.66-71.58]
Onset site	Bulbar: 449 (31%) Axial: 3 (0.002%) Generalized: 4 (0.003%) Limbs: 998 (68%)	Bulbar: 105 (30%) Axial: 0 (0%) Generalized: 0 (0%) Limbs: 242 (70%)
ALSFRS-R slope	0.43 [0.24-0.79]	0.41 [0.23-0.80]
Number of available visits	2.00 [2.00-3.00]	3.00 [2.00-3.00]

Runs should be uploaded using the following name convention for their identifiers:

`<teamname>_T<1|2><a|b|c>_<train>_<freefield>`

where:

- `teamname` is the name of the participating team;
- `T<1|2><a|b|c>` is the identifier of the task the run is submitted to, e.g. T1b for Task 1, sub-task b;
- `train` is data window used to train the algorithm. It can be either `M0`, if only the data available at Time 0 have been used, or `M6` if all the data available in the first 6 months have been used;
- `freefield` is a free field that participants can use as they prefer.

For example, a complete run identifier may look like

`upd_T2b_M6_survRF`

where:

- `upd` is the University of Padua team;
- `T2b` means that the run is submitted for Task 2, sub-task b;
- `M6` means that the algorithm has been trained using all the data available in the first 6 months;
- `survRF` suggests that participants have used survival random forests as a prediction method.

Participant Repository Participants are provided with an individual git repository for all the tasks they take part in. The repository contains the runs, resources, and possibly the code produced by each participant in order to promote reproducibility and open science. The repository is organised as follows:

- submission: this folder contains the runs submitted for the different tasks.
- score: this folder contains the performance scores of the submitted runs.
- code: this folder contains the source code of the developed system.
- resource: this folder contains any additional resources created during the participation.
- report: this folder contains the template for participant report.

The submission and score folders are organized into sub-folders for each task as follows:

- submission/task1: for the runs submitted to the first task. Similar structure for the other tasks.
- score/task1: for the performance scores of the runs submitted to the first task. Similar structure for the other tasks.

The goal of iDPP:CLEF is to speed up the creation of systems and resources for ALS progression prediction as well as openly share these systems and resources as much as possible. Therefore, participants are more than encouraged to share their code and any additional resources they have used or created.

All the contents of these repositories are released under the *Creative Commons Attribution-ShareAlike 4.0 International License*[11].

Task 1 Run Format Runs had to be submitted as a text file with the following format:

```
0x4bed50627d141453da7499a7f6ae84ab 0.897 0 PEG upd_T1b_M6_survRF
0x4d0e8370abe97d0fdedbded6787ebcfc 0.773 1 PEG upd_T1b_M6_survRF
0x5bbf2927feefd8617b58b5005f75fc0d 0.773 2 DEATH upd_T1b_M6_survRF
0x814ec836b32264453c04bb989f7825d4 0.615 3 NONE upd_T1b_M6_survRF
0x71dabb094f55fab5fc719e348dffc85 0.317 4 PEG upd_T1b_M6_survRF
...
```

where:

- Columns are separated by a white space;
- The first column is the `patient ID`, a 128 bit hex number (should be considered just as a string);
- The second column shows the prediction score that generated the ranking. It is expected to be a floating point number in the range $[0, 1]$. This score must be in descending (non-increasing) order;

[11] http://creativecommons.org/licenses/by-sa/4.0/

- The third column is the rank of the patient by her/his risk of impairment, starting from 0. This is expected to be a strictly increasing integer number. It is important to include the rank so that we can handle tied scores (for a given run) in a uniform fashion;
- The fourth column is the predicted event. It comes from a controlled vocabulary and it can be either `NIV`, `PEG`, `DEATH`, or `NONE`. Note that, since each sub-task is focused on the prediction of a specific event (`NIV`, `PEG`, or `DEATH`), this column will contain that event or the competing event `DEATH` or `NONE`;
- The fifth column is the run identifier, according to the format described above. It must uniquely identify the participating team and the submitted run.

Task 2 Run Format Runs had to be submitted as a text file with the following format:

```
0x4bed50627d141453da7499a7f6ae84ab 6-12 PEG upd_T2b_M6_survRF
0x4d0e8370abe97d0fdedbded6787ebcfc 18-24 PEG upd_T2b_M6_survRF
0x5bbf2927feefd8617b58b5005f75fc0d 24-30 DEATH upd_T2b_M6_survRF
0x814ec836b32264453c04bb989f7825d4 >36 NONE upd_T2b_M6_survRF
0x71dabb094f55fab5fc719e348dffc85 >36 PEG upd_T2b_M6_survRF
...
```

where:

- Columns are separated by a white space;
- The first column is the `patient ID`, a 128 bit hex number (should be considered just as a string);
- The second column shows the prediction window in months. Possible values are taken from a controlled vocabulary as follows:
  - `6-12`: the event will happen in the range of months $(6, 12]$;
  - `12-18`: the event will happen in the range of months $(12, 18]$;
  - `18-24`: the event will happen in the range of months $(18, 24]$;
  - `24-30`: the event will happen in the range of months $(24, 30]$;
  - `30-36`: the event will happen in the range of months $(30, 36]$;
  - `>36`: the event will happen in the range of months $(36, +\infty)$.
- The third column is the rank of the patient by her/his risk of impairment, starting from 0. It is important to include the rank so that we can handle tied scores (for a given run) in a uniform fashion;
- The fourth column is the predicted event. It comes from a controlled vocabulary and it can be either `NIV`, `PEG`, `DEATH`, or `NONE`. Note that, since each sub-task is focused on the prediction of a specific event (`NIV`, `PEG`, or `DEATH`), this column will contain that event or the competing event `DEATH` or `NONE`;
- The fifth column is the run identifier, according to the format described above. It must uniquely identify the participating team and the submitted run.

Table 4: Teams participating in iDPP:CLEF 2022.

Team Name	Description	Country	Repository	Paper
BioHIT	National Centre for Scientific Research Demokritos (NCSR Demokritos)	Greece	https://bitbucket.org/brainteaser-health/idpp2022-biohit	–
CompBioMed	Department of Medical Sciences, University of Turin	Italy	https://bitbucket.org/brainteaser-health/idpp2022-compbiomed-unito	Pancotti et al. [16]
FCOOL	Faculty of Sciences of the University of Lisbon	Portugal	https://bitbucket.org/brainteaser-health/idpp2022-fcool	Branco et al. [2] and Nunes et al. [15]
LIG GETALP	Laboratoire d'Informatique de Grenoble, Université Grenoble Alpes	France	https://bitbucket.org/brainteaser-health/idpp2022-lig-getalp	Mannion et al. [14]
SBB	University of Padua	Italy	https://bitbucket.org/brainteaser-health/idpp2022-sbb	Trescato et al. [18]

Table 5: Break-down of the runs submitted by participants for each task and sub-task. Participation in Task 3 does not involve submission of runs and it is marked just with a tick.

Team Name	Total	Task 1			Task 2			Task 3
		a	b	c	a	b	c	
BioHIT	18	3	3	3	3	3	3	–
CompBioMed	40	8	8	6	6	6	6	–
FCOOL	15	–	–	–	5	5	5	✓
LIG GETALP	23	4	4	4	4	4	3	–
SBB	24	4	4	4	4	4	4	–
Total	**120**	**19**	**19**	**17**	**22**	**22**	**21**	

5.2 Participants

Overall, 43 teams registered for participating in iDPP:CLEF but only 5 of them actually managed to submit runs for at least one of the offered tasks. Table 4 reports the details about the participating teams.

Table 5 provides breakdown of the number of runs submitted by each participant for each task and sub-task. Overall, we have received 120 runs which are roughly broken down evenly among the different tasks.

6 Evaluation Measures

iDPP:CLEF adopted several state-of-the-art evaluation measures to assess the performance of the prediction algorithms, among which:

– *ROC curve and/or the precision-recall curve (and area under the curve)* to show the trade-off between clinical sensitivity and specificity for every possible cut-off of the risk scores;

- *Concordance Index (C-index)* to summarize how well a predicted risk score describes an observed sequence of events.
- *E/O ratio and Brier Score* to assess whether or not the observed event rates match expected event rates in subgroups of the model population.
- *Specificity and recall* to assess, for each interval, the ability of the models of correctly identify true positives and true negatives.
- *Distance* to assess how far the predicted time interval was from the true time interval.

To ease the computation and reproducibility of the results, scripts for computing the measures are available in the following repository: https://bitbucket.org/brainteaser-health/idpp2022-performance-computation.

The next two sections provide details about the adopted measures for each Task.

6.1 Pilot Task 1: Ranking Risk of Impairment

The runs submitted for Task 1 were evaluated by means of Harrel's concordance index (C-index) [11], area under the receiver operating characteristic curve (AUROC) [10], and the Brier score (BS) [3]. The 95% confidence intervals of the C-index and the AUROC were also considered [17].

The C-index has an advantage over the other considered metrics (i.e., AUROC and BS) in that it can be used to evaluate model discrimination on the test sets regardless of censored data – data for those patients that did not incurred either the relevant event (NIV or PEG) or the competing event (Death). According to the best practices in the field [13], before computing the C-index, a final censoring time equal to the last time-to-event in the training was set on each test set. This ensured consistency between Task 1's final results and those that might have been obtained by the participants during model development.

The AUROC and BS were computed at various prediction horizons (PHs). Specifically, seven clinically relevant PHs were considered, namely: 12, 18, 24, 30, 36, 48, and 60 months after the baseline. For each PH, the corresponding version of the test set comprised: all patients who experienced an event before the PH, and all patients who experienced an event or were censored after the PH as censored patients (and were, thus, censored at that PH). As the status of patients censored before the PH was, by definition, unknown, they were excluded from performance evaluation at that PH.

To contextualize the results obtained by the participants, each run was compared to the empirical lower bound established by the average performance of 100 random classifiers (i.e., such that their output was a random continuous number, uniformly sampled in the range $[0, 1]$).

6.2 Pilot Task 2: Predicting Time of Impairment

To evaluate the predictions of Task 2, the selected evaluation metrics were: the specificity, the recall, and a measure of distance between the predicted and correct time intervals.

Confusion matrices were computed to derive specificity, i.e., the number of correct negative predictions divided by the total number of negatives, and recall, i.e., the ratio of correct positive predictions over the total predicted positives. To do so, the outcome times reported in the column *Time* of the published test sets were mapped to the corresponding interval ("6-12", "12-18", "18-24", "24-30", "30-36", or ">36" months). A conformance check was performed on the participants' predicted times: predictions in the time interval "0-6" were reassigned to the interval "6-12", i.e., the closest allowed interval. The confusion matrices reported the predicted time interval vs the true time interval, independently of the predicted event.

A measure of distance between the predicted and correct time intervals, in months, was also considered (AbsDist). To compute the AbsDist, all the time intervals were replaced with the mean value of each interval (i.e., "6-12" was replaced with 9, "12-18" with 15, "18-24" with 21, "24-30" with 27, "30-36" with 33, and ">36" with 39). The difference between the predicted values and the true values was then computed as $meanValue_{predicted\,time\,interval} - meanValue_{true\,time\,interval}$. The obtained differences were, by construction, in the range $[-36; +36]$ where a smaller modulus corresponds to more accurate predictions. Negative values correspond to a events that occur before the predicted time and positive values to events that occur after. Finally, the AbsDist was obtained by averaging the differences absolute values.

To contextualize the results obtained by the participants, each run was compared to the performance of several synthetic runs, with the following characteristics:

- *min_interval*: a run in which the predicted time intervals are identical for all subjects, and fixed at the first possible time interval, i.e. "6-12";
- *max_interval*: a run in which the predicted time intervals are identical for all subjects, and fixed at the last possible time interval, i.e. ">36";
- *interval_18_24*: a run in which the predicted time intervals are identical for all subjects, and fixed at the time interval "18-24";
- *random_interval*: 100 randomly generated runs, but with the same distribution as the test set distribution (i.e., such that their output was sampled among the labels "6-12", "12-18", "18-24", "24-30", "30-36", ">36" following the same distribution of the true intervals);
- *inverse_distr_interval*: 100 randomly generated runs, but with an inverse distribution compared to the test set distribution (i.e., such that their output was sampled among the labels "6-12", "12-18", "18-24", "24-30", "30-36", ">36" following the inverse distribution of the true outcome);
- *corr_interval*: 100 correlated runs, with correlation coefficient to the true intervals ~ 0.7.

7 Results

For each task, we report here the analysis of the performance attained by the runs submitted by the Lab's participants according to the metrics described in Section 6.

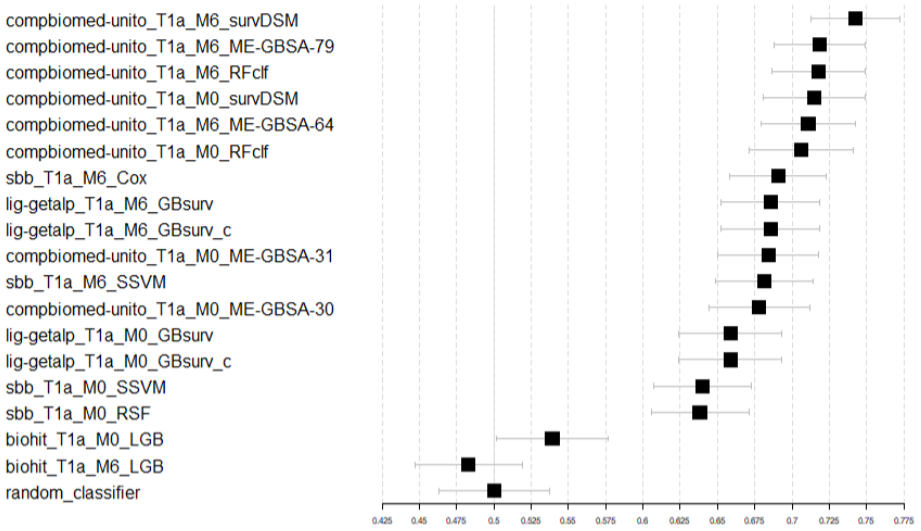

Fig. 4: Sub-task a C-index computed for all submitted runs. The bars in the plot show the 95% confidence intervals. The average C-index of 100 random classifiers is reported in the last row.

7.1 Pilot Task 1: Ranking Risk of Impairment for ALS

Here, only the C-index, the 12-month and 48-month AUROCs, and the 12-month and 48-month BSs obtained for all participants runs submitted for Task 1's sub-task a are shown. Results for all sub-tasks and all PHs are shown in the extended overview [9].

Figure 4 shows the C-index with its 95% confidence intervals computed for all runs submitted for sub-task a and for the 100 random classifiers (last row). As expected, the random classifiers yielded an average C-index of around 0.5. Runs submitted by the BioHit team were comparable to those obtained by the random classifiers. All runs submitted by other participants significantly outperformed the random classifiers (C-index > 0.625) with team CampBioMed leading the pack (C-index > 0.7).

Figure 5 shows the AUROC with its 95% confidence intervals computed for all runs submitted for sub-task a at the 12-month PH. The average 12-month AUROC of the 100 random classifiers is reported in the the last row. The 12-month AUROC confirmed the results obtained when considering the C-index. Again, as expected, the random classifiers yielded a 12-month AUROC of around 0.5. Runs submitted by the BioHit team showed a discrimination that was comparable to the one of the random classifiers, and all runs submitted by other participants significantly outperformed the random classifiers (12-month AUROC > 0.675) with some runs of team CampBioMed and team SBB achieving

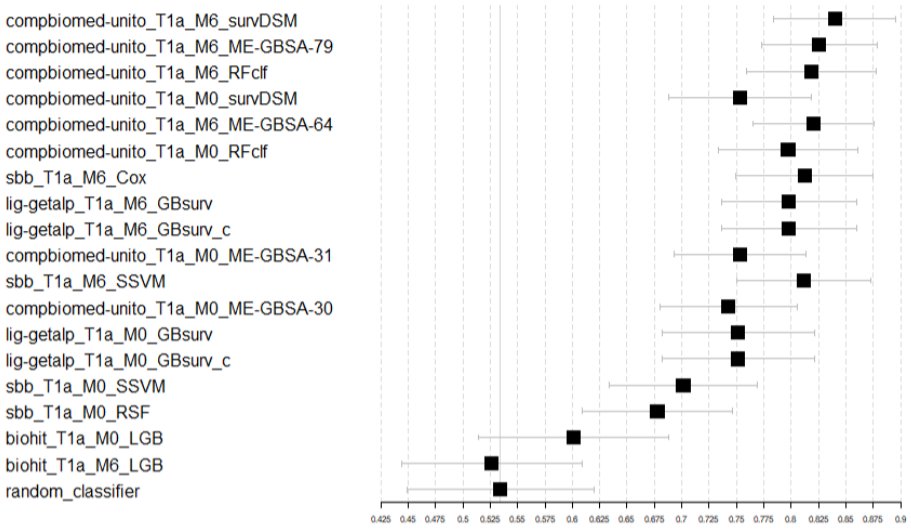

Fig. 5: Sub-task a AUROC computed for all submitted runs with a 12-months PH. The bars in the plot show the 95% confidence intervals. The average 12-months AUROC of 100 random classifiers is reported in the last row.

excellent results (12-month AUROC > 0.8) when all the information available in the first 6 months was considered (M6 runs).

Figure 6 shows the BS computed for runs submitted for sub-task a at the 12-month PH. The average 12-month BS of the 100 random classifiers is reported in the last row. The random classifier yield a 12-month BS of around 0.325 as the random probability values were, on average, well distributed in the range $[0, 1]$. Runs submitted by the CampBioMed team showed the best calibration at this PH (12-month BS < 0.225), while those submitted by the SBB team showed the worst one (12-month BS > 0.675), mainly due to a consistent overestimation of the event probability. Other participants' runs had 12-month BSs comparable with the random classifiers as their models did not correctly predict the event probability but neither showed consistent overestimation trends.

Figure 7 shows the AUROC with its 95% confidence intervals computed for all runs submitted for sub-task a at the 48-month PHs. The average 48-month AUROC of the 100 random classifiers (again, expectedly, around 0.5) is reported in the the last row. The 48-month AUROC confirmed once again the results obtained with the C-index and 12-month AUROC. Runs submitted by the BioHit team had comparable discrimination to the random classifiers, while all runs submitted by other participants significantly outperformed them (48-month AUROC > 0.7). Runs that used all the information available in the first 6 months (M6 runs) submitted by the CampBioMed team were the best performing ones also at this PH (48-month AUROC > 0.8).

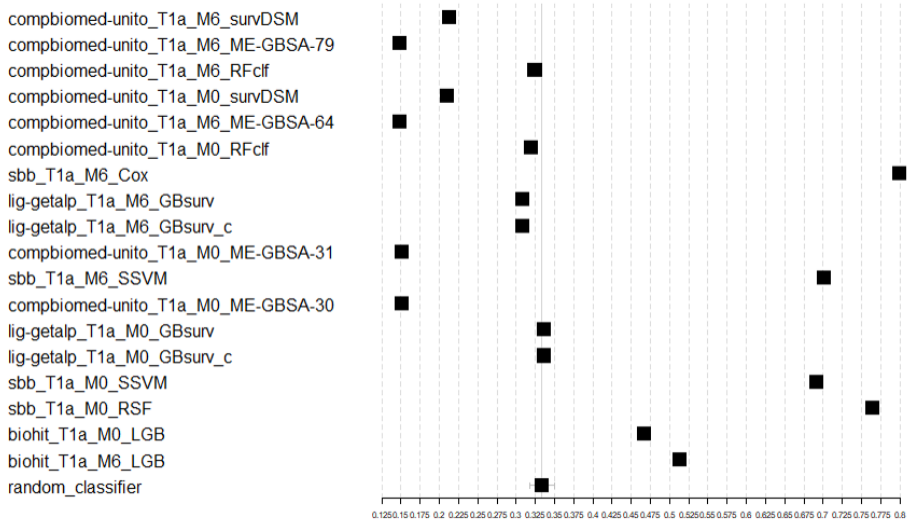

Fig. 6: Sub-task a BS computed for all submitted runs with a 12-months PH. The random classifier average 12-months BS is reported in the last row with its 95% confidence intervals.

Figure 8 shows the BS computed for runs submitted for sub-task a at the 48-month PH. The average 48-month BS of the 100 random classifiers is reported in the last row. The random classifiers yielded a 48-month BS of around 0.325 as the random probability values were, as in the 12-month case, on average, well distributed in the range $[0, 1]$. All GBSA runs submitted by the CampBioMed team, which had good calibration with a PH of 12 months, led to a poorer calibration at 48 months (48-month BS > 0.75). All other runs submitted by the participants significantly outperformed the random classifiers by showing good calibration at this PH (48-months BS < 0.25).

Overall, for Task 1 sub-task a, runs submitted by the CampBioMed team were the best performing across the board; meanwhile, runs submitted by the BioHit team led to the lowest discrimination, but still yielded acceptable calibration at a long PH (48 months). Finally, the SBB and LIG GETALP teams obtained comparable results when considering runs obtained using all the information available in the first 6 months (M6 runs); meanwhile, when using only the information available at time 0 (M0 runs), runs submitted by the SBB team showed worse discrimination than those submitted by the LIG GETALP teams.

7.2 Pilot Task 2: Predicting Time of Impairment for ALS

Figures 9, 10, and 11 show the specificity-recall plots for three select time intervals ("6-12", "12-18", and "18-24") of Task 2's sub-task a, including all participants' runs and all the synthetic runs. Results for all time intervals and sub-tasks

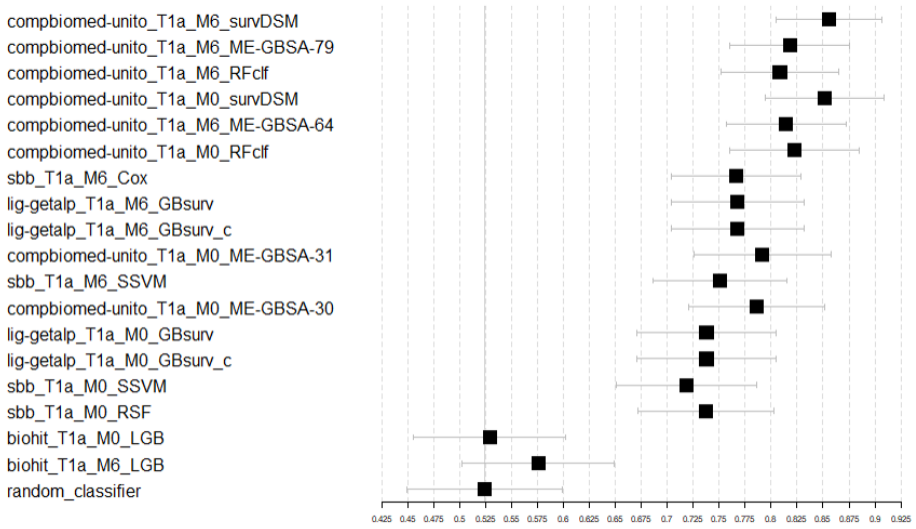

Fig. 7: Sub-task a AUROC computed for all submitted runs with a 48-months PH. The bars in the plot show the 95% confidence intervals. The average 48-months AUROC of 100 random classifiers is reported in the last row.

are presented in the extended overview [9]. The graph shows the specificity on the x-axis (from 1 to 0, left to right), and the recall on the y-axis (from 0 to 1, bottom to top). The ideal classifier would have specificity = 1 and recall = 1, and would therefore be located in the upper left corner: as a general guidance, the closer a run is to the upper left corner, the better the classification obtained.

In all graphs, the synthetic runs with constant predictions, fixed at the minimum or maximum allowed interval, are located in the two extreme corners of the plot. In detail, the *max_interval* run is located in the lower left corner with specificity = 0 and recall = 1, while the *min_interval* run, in the upper right corner, has specificity = 1 and recall = 0. As expected, the 100 runs with 70% correlation form a cloud in the upper left corner, while the 200 randomly generated runs, 100 with the same distribution and 100 with the inverse distribution always remain in the lower left sector, with $1 > specificity > 0.5$ and $0 > recall > 0.5$.

For the "6-12" interval, represented in Figure 9, the team with the best classification performance according to specificity and recall was the FCOOL team, whose five submitted runs yielded specificity $\simeq 0.72$ and recall = 0.612. One run from the CompBioMed team also performed well, with specificity = 0.839 and recall = 0.561. In contrast, the other runs submitted by the CompBioMed team were in line with those of the other participants, with rather high specificity but low recall.

Figure 10 reports the results for the "12-18" interval. Again, the CompBioMed team outperformed the other teams with a run with specificity = 0.581 and recall = 0.545. The second best run in this time window was from the LIG

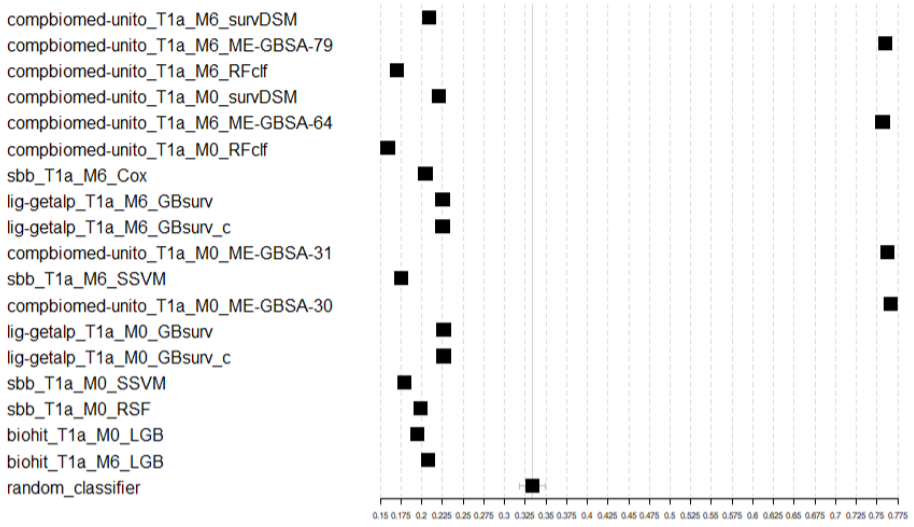

Fig. 8: Sub-task a BS computed for all submitted runs with a 48-months PH. The random classifier average 48-months BS is reported in the last row with its 95% confidence intervals.

GETALP team, with specificity = 0.668 and recall = 0.509. Similar results were obtained by two runs of the SBB team, which reached specificity $\simeq$ 0.64 and recall $\simeq$ 0.44.

The SBB and LIG GETALP teams obtained the best performance in the interval "18-24", as shown in Figure 11. Specifically, the SBB team submitted three runs with specificity $\simeq$ 0.66 and recall $\simeq$ 0.47, while LIG GETALP one with specificity = 0.697 and recall = 0.479. The other teams, as in the other time intervals, obtained higher specificity scores to the detriment of recall, thus positioning themselves in the lower left quadrant.

Figure 12 shows the AbsDist computed for all runs submitted for sub-task a. The average AbsDist of the synthetic runs is reported as well. As expected, the *max_interval* run led to the worst result (AbsDist $>$ 17 months), as most subjects have a true time interval smaller than the maximum one. Runs *random_interval*, *min_interval*, and *inverse_corr_interval* led to comparable distance values (AbsDist $12-13$ months). Runs submitted by the BioHit team had AbsDist values comparable with the synthetic run *interval_18_24* (AbsDist $10-11$ months), suggesting that their models might predict the average time interval for most subjects. All runs submitted by the other teams significantly outperformed the aforementioned synthetic runs (AbsDist $7-9$ months) with CampBioMed team leading the pack. Finally, the *corr_interval* run led to the smallest AbsDist value (AbsDist $<$ 4 months). Note, however, that this run was included only as an arbitrary reference, and its distance value was not strictly expected to be reached by any participant.

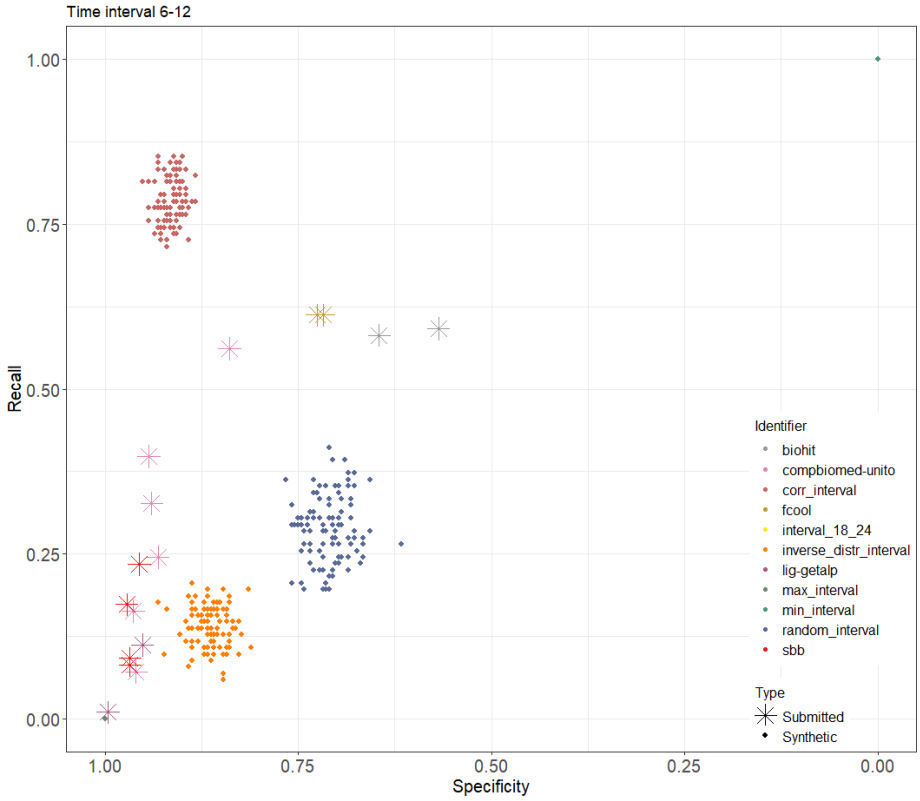

Fig. 9: Specificity-recall plot, sub-task a, time interval 6-12.

Predicting the correct event time interval proved to be a challenge for all teams, especially in terms of recall. However, almost all teams were able to obtain good AbsDist values as, on average, their models, despite not being able to precisely identify the correct time interval, tended to predict an interval that was immediately before or after the true one.

As observed for Task 1, runs performed better when considering all the information available in the first 6 months (M6 runs) rather than only the information available at time 0 (M0 runs).

7.3 Approaches

In this section, we provide a short summary of the approaches adopted by participants in iDPP@CLEF. There are two separate sub-sections, one for Task 1 and 2 focused on ALS progression prediction and the other for Task 3, on *eXplainable AI (XAI)* approaches for such kind of algorithms.

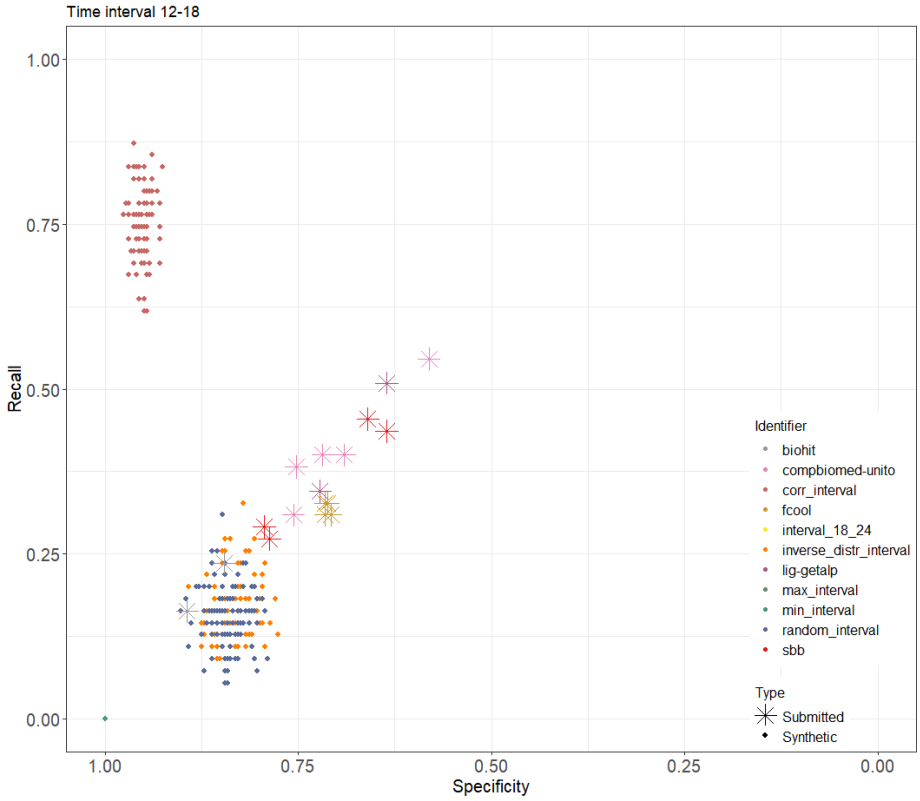

Fig. 10: Specificity-recall plot, sub-task a, time interval 12-18.

Task 1 and 2 BoiHIT explored the use of logistic regression, random forest classifiers, XGBoost, and LightGBM. Decision trees and boosting approaches were preferred due to their ability to deal with both categorical and numerical/continuous features and the interpretability they offer. Even if LightGBM was the model with the best performance, BoiHIT found out that this kind of approaches might not be appropriate for time dependent problems and that time to event analysis methods, such as survival analysis, might yield better results.

CompBioMed [16] considered three main approaches. The simplest one consisted on fitting a standard survival predictor separately for each event as outlined above for independent events, called Naive Multiple Event Survival (NMES). Another was the recently developed Deep Survival Machine (DSM), based on deep learning and capable of handling competing risks. Finally, they also proposed a time-aware classifier ensemble method, that also handles competing risks, called Time-Aware Classifier Ensemble (TACE). All the above approaches achieved comparable performance among them. Only the TACE models appeared to be slightly worse than the rest in when using 6 months of data. Moreover, no clear advantage of the DSM models, that specifically handles competing

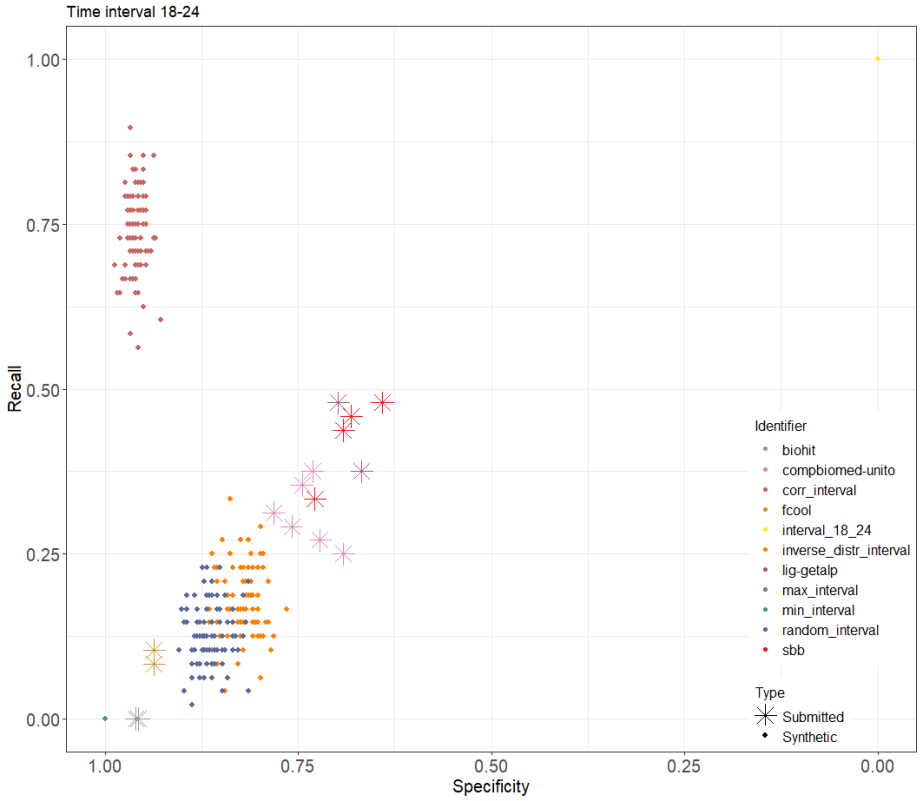

Fig. 11: Specificity-recall plot, sub-task a, time interval 18-24.

risks, was observed with respect to the NMES models, which treat all events, as if they were independent.

FCOOL [2] proposes a hierarchical approach, with a first-stage event prediction, followed by specialized models predicting the time window to a particular event. The procedure is three-fold: first, it creates patient snapshots based on clustering with constraints, thus organizing patient records in an efficient manner. Second, it uses a pattern-based approach that incorporates recent advances on temporal pattern mining to the context of classification. This approach performs end-stage event prediction while allowing the entire patient's medical history to be considered. Finally, exploiting the predictions from the previous step, specialized models are learned using the original features to predict the time window to an event. This two-stage prediction approach aimed to promote homogeneity and lessen the impact of class imbalance, in comparison to performing one single multilabel task.

LIG GETALP [14] employed Cox's proportional hazards model to the task of ranking the risk of impairment, using the gradient boosting learning strategy The output of the time-independent part of the survival function calculated by

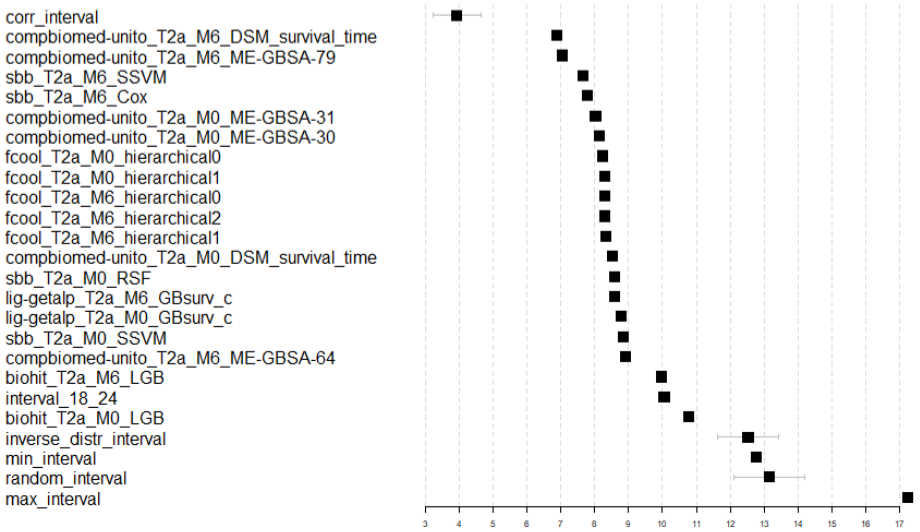

Fig. 12: Sub-task a AbsDist computed for all submitted and synthetic runs. The AbsDist of *corr_interval*, *inverse_distr_interval*, and *random_interval* is the average with 95% confidence intervals computed on the corresponding 100 randomly generated runs.

the gradient boosting survival analysis method is then mapped to the interval (0, 1), via a sigmoid function. To estimate the time-to-event, LIG GETALP used a regression model based on Accelerated Gradient Boosting (AGB). This being a standard regression model, it does not take censoring into account and Mannion et al. uses class predictions based on the Task 1 survival model to "censor" the time-to-event predictions.

SBB [18] considered three survival analysis methods, namely: Cox, SSVM, and RSF. They were chosen to represent a broad spectrum of baseline models including parametric (SSVM), semiparametric (Cox), linear (Cox, SSVM), and nonlinear (RSF) models. The Cox model and the RSF can only output risk scores, which can be used to address Task 1 by ranking ALS patients according to their risk of impairment, but do not provide a straightforward solution to predicting Task 2's time of impairment. To extend these approaches to Task 2, the predicted time of impairment for a given patient was selected as the median predicted time to impairment, i.e., the time at which the estimated survival function crossed the 0.5 threshold. Instead, the SSVM can be used either as a ranker or a time regressor depending on how the risk ratio hyperparameter is set during model training. Here, the SSVM was initially trained as a time regressor to address Task 2 directly. Then, its predicted times were converted into risk scores in the range [0-1], as requested by the challenge rules, via Platt scaling.

Task 3 Nunes et al. [15] proposes a novel approach that generates semantic similarity-based explanations for patient-level predictions. The underlying idea is to explain the prediction for one patient by considering aspect-oriented semantic similarity with other relevant patients based on the most important features used by ML approaches or selected by users. To build rich and easy to understand semantic-similarity based explanations, Nunes et al. developed five steps: (1) the enrichment of the Brainteaser Ontology [1] through integration of other biomedical ontologies; (2) the semantic annotation of patients (if not already available); (3) the similarity calculation between patients; (4) selection of the set of patients to explain a specific prediction; and (5) the visualization of the generated similarity-based explanations.

Buonocore et al. [4] trained a set of 4 well-known classifiers to predict death occurrence: Gradient Boosting (using XGB implementation), Random Forest, Logistic Regression and Multilayer perceptron. For the XAI methods Buonocore et al. focused our attention on three different methods for post-hoc, model-agnostic, local explainability, selecting SHAP, LIME and AraucanaXAI. Then, Buonocore et al. evaluated and compared XAI approaches in terms of a set of metrics defined in previous research on XAI in healthcare: *identity*: if there are two identical instances, they must have the same explanations; *fidelity*: concordance of the predictions between the XAI surrogate model and the original ML model; *separability*: if there are 2 dissimilar instances, they must have dissimilar explanations; *time*: average time required by the XAI method to output an explanation across the entire test set. The quantitative evaluation of the three different XAI methods did not reveal definitive superior performance of one of the approaches, albeit SHAP seems to be the better overall performing algorithm. However the explainability evaluation metrics are not all that is needed to thoroughly assess the multifaceted construct of what constitutes a "good" explanation in XAI in healthcare.

8 Conclusions and Future Work

iDPP:CLEF is a new pilot activity focusing on predicting the temporal progression of ALS and on the explainability of the AI algorithms for such prediction.

We developed 3 datasets containing anonymized patient data from two medical institutions, one in Turin and the other in Lisbon, for the prediction of NIV, PEG, or death.

Out of 43 registered participants, 5 managed to submit a total of 120 runs, evenly spread across the offered tasks. Participants adopted a range of approaches, including various types of survival analysis, also using deep learning techniques. For the XAI of the prediction algorithms they used both semantic-similarity based techniques and state-of-art post-hoc and model-agnostic XAI approaches.

For this initial iteration of the lab, iDPP:CLEF focus on ALS progression prediction. Possible, future cycles will be extended to *Multiple Sclerosis (MS)*, another chronic disease, impairing neurological functions. Moreover, we plan to

extend the datasets to also include data from environmental sensor, e.g. concerning pollution.

Acknowledgments

The work reported in this paper has been partially supported by the BRAINTEASER[12] project (contract n. GA101017598), as a part of the European Union's Horizon 2020 research and innovation programme.

References

1. Bettin, M., Di Nunzio, G.M., Dosso, D., Faggioli, G., Ferro, N., Marchetti, N., Silvello, G.: Deliverable 9.1 – Project ontology and terminology, including data mapper and RDF graph builder. BRAINTEASER, EU Horizon 2020, Contract N. GA101017598. https://brainteaser.health/ (December 2021)
2. Branco, R., Soares, D., Martins, A.S., Auletta, E., Castanho, E.N., Nunes, S., Serrano, F., Sousa, R.T., Pesquita, C., Madeira, S.C., Aidos, H.: Hierarchical Modelling for ALS Prognosis: Predicting the Progression Towards Critical Events. In: [7]
3. Brier, G.W.: Verification of Forecasts Expressed in Terms of Probability. Monthly Weather Review **78**(1), 1–3 (January 1950)
4. Buonocore, T.M., Nicora, G., Dagliati, A., Parimbelli, E.: Evaluation of XAI on ALS 6-months mortality prediction. In: [7]
5. Cedarbaum, J.M., Stambler, N., Malta, E., Fuller, C., Hilt, D., Thurmond, B., Nakanishi, A.: The ALSFRS-R: a revised ALS functional rating scale that incorporates assessments of respiratory function. Journal of the Neurological Sciences **169**(1–2), 13–21 (October 1999)
6. Chio, A., Logroscino, G., Hardiman, O., Swingler, R., Mitchell, D., Beghi, E., Traynor, B.G., Consortium, E., et al.: Prognostic factors in als: a critical review. Amyotrophic lateral sclerosis **10**(5-6), 310–323 (2009)
7. Faggioli, G., Ferro, N., Hanbury, A., Potthast, M. (eds.): CLEF 2022 Working Notes, CEUR Workshop Proceedings (CEUR-WS.org), ISSN 1613-0073 (2022)
8. Guazzo, A., Trescato, I., Longato, E., Hazizaj, E., Dosso, D., Faggioli, G., Di Nunzio, G.M., Silvello, G., Vettoretti, M., Tavazzi, E., Roversi, C., Fariselli, P., Madeira, S.C., de Carvalho, M., Gromicho, M., Chiò, A., Manera, U., Dagliati, A., Birolo, G., Aidos, H., Di Camillo, B., Ferro, N.: Intelligent Disease Progression Prediction: Overview of iDPP@CLEF 2022. In: Barrón-Cedeño, A., Da San Martino, G., Degli Esposti, M., Sebastiani, F., Macdonald, C., Pasi, G., Hanbury, A., Potthast, M., Faggioli, G., Ferro, N. (eds.) Experimental IR Meets Multilinguality, Multimodality, and Interaction. Proceedings of the Thirteenth International Conference of the CLEF Association (CLEF 2022), Lecture Notes in Computer Science (LNCS) 13390, Springer, Heidelberg, Germany (2022)

[12] https://brainteaser.health/

9. Guazzo, A., Trescato, I., Longato, E., Hazizaj, E., Dosso, D., Faggioli, G., Di Nunzio, G.M., Silvello, G., Vettoretti, M., Tavazzi, E., Roversi, C., Fariselli, P., Madeira, S.C., de Carvalho, M., Gromicho, M., Chiò, A., Manera, U., Dagliati, A., Birolo, G., Aidos, H., Di Camillo, B., Ferro, N.: Overview of iDPP@CLEF 2022: The Intelligent Disease Progression Prediction Challenge. In: [7]
10. Hanley, J.A., McNeil, B.J.: The meaning and use of the area under a receiver operating characteristic (roc) curve. Radiology **143**(1), 29–36 (1982), pMID: 7063747
11. Harrell, Frank E., J., Califf, R.M., Pryor, D.B., Lee, K.L., Rosati, R.A.: Evaluating the Yield of Medical Tests. JAMA **247**(18), 2543–2546 (05 1982), ISSN 0098-7484
12. Küffner, R., Zach, N., Norel, R., Hawe, J., Schoenfeld, D., Wang, L., Li, G., Fang, L., Mackey, L., Hardiman, O., Cudkowicz, M., Sherman, A., Ertaylan, G., Grosse-Wentrup, M., Hothorn, T., van Ligtenberg, J., Macke, J.H., Meyer, T., Schölkopf, B., Tran, L., Vaughan, R., Stolovitzky, G., Leitner, M.L.: Crowdsourced analysis of clinical trial data to predict amyotrophic lateral sclerosis progression. Nature Biotechnology **33**(1), 51–57 (January 2015)
13. Longato, E., Vettoretti, M., Di Camillo, B.: A practical perspective on the concordance index for the evaluation and selection of prognostic time-to-event models. Journal of Biomedical Informatics **108**, 103496:1–103496:9 (August 2020)
14. Mannion, A., Chevalier, T., Schwab, D., Goeuriot, L.: Predicting the Risk of & Time to Impairment for ALS patients. In: [7]
15. Nunes, S., Sousa, R.T., Serrano, F., Branco, R., Soares, D.F., Martins, A.S., Auletta, E., Castanho, E.N., Madeira, S.C., Aidos, H., Pesquita, C.: Explaining Artificial Intelligence Predictions of Disease Progression with Semantic Similarity. In: [7]
16. Pancotti, C., Birolo, G., Sanavia, T., Rollo, C., Fariselli, P.: Multi-Event Survival Prediction for Amyotrophic Lateral Sclerosis. In: [7]
17. Pencina, M.J., D'Agostino, R.B.: Overall c as a measure of discrimination in survival analysis: model specific population value and confidence interval estimation. Statistics in Medicine **23**(13), 2109–2123 (2004)
18. Trescato, I., Guazzo, A., Longato, E., Hazizaj, E., Roversi, C., Tavazzi, E., Vettoretti, M., Di Camillo, B.: Baseline Machine Learning Approaches To Predict Amyotrophic Lateral Sclerosis Disease Progression. In: [7]