



OPEN

# Maximally informative feature selection using Information Imbalance: Application to COVID-19 severity prediction

Romina Wild<sup>1</sup>, Emanuela Sozio<sup>2,3</sup>, Riccardo G. Margiotta<sup>1</sup>, Fabiana Dellai<sup>2</sup>, Angela Acquasanta<sup>2</sup>, Fabio Del Ben<sup>3</sup>, Carlo Tascini<sup>2,3</sup>, Francesco Curcio<sup>3,2</sup> & Alessandro Laio<sup>1,4</sup>✉

Clinical databases typically include, for each patient, many heterogeneous features, for example blood exams, the clinical history before the onset of the disease, the evolution of the symptoms, the results of imaging exams, and many others. We here propose to exploit a recently developed statistical approach, the Information Imbalance, to compare different subsets of patient features and automatically select the set of features that is maximally informative for a given clinical purpose, especially in minority classes. We adapt the Information Imbalance approach to work in a clinical framework, where patient features are often categorical and are generally available only for a fraction of the patients. We apply this algorithm to a data set of ~1300 patients treated for COVID-19 in Udine hospital before October 2021. Using this approach, we find combinations of features which, if used in combination, are maximally informative of the clinical fate and of the severity of the disease. The optimal number of features, which is determined automatically, turns out to be between 10 and 15. These features can be measured at admission. The approach can be used also if the features are available only for a fraction of the patients, does not require imputation and, importantly, is able to automatically select features with small inter-feature correlation. Clinical insights deriving from this study are also discussed.

In many fields, and in particular in statistical medicine, one attempts to develop a predictor using relatively few data points (the patients), characterized by a number of features which can be large. These features encompass demographics, vital parameters, comorbidities, medications, blood test values, radiological exams, clinical scores and more. Furthermore, they can be of any data type, e.g. quantitative (weight, age, blood value levels), binary (presence of diabetes or other comorbidities), nominal (types of ventilation), or ordinal (sequential organ failure assessment (SOFA) score). Many of these features are typically irrelevant or redundant, namely correlated with each other, and a selection of few, relevant features is desirable. For medical professionals, having to consider too many features confuses the clinical work.

Typically, a feature is considered relevant if it correlates with the target, for example if it discriminates between target classes. This simple concept is at the basis of most feature selection algorithms. Feature selection methods can be broadly divided into filter, wrapper and embedded methods. Filters are simple statistics to rank the features independently of the subsequent prediction (classifier agnostic), while wrapper and embedded methods use the predictor as criterion to select feature subsets<sup>1,2</sup>. Among the classic embedded feature selection methods, lasso-regularized regression<sup>3</sup> provides interpretable feature selection while drawing a linear relationship between input features and target. Sparse additive models (SPAMs)<sup>4</sup> and sparse neural additive models (SNAMs)<sup>5</sup> extend this to the non-linear case. Variable ranking filter methods tend to be easier and faster to use than other feature selection methods, but they have the drawback of being univariate methods with inability to find the optimal dimension of the feature space, namely the number of variables that are necessary to make a good prediction<sup>1,6</sup>. Furthermore, many existing feature selection algorithms suffer from the inability to handle missing and noisy data<sup>2</sup>.

<sup>1</sup>International School for Advanced Studies (SISSA), Via Bonomea 265, 34136 Trieste, Italy. <sup>2</sup>Infectious Disease Unit, Azienda Sanitaria Universitaria Friuli Centrale (ASU FC), Via Pozzuolo 330, 33100 Udine, Italy. <sup>3</sup>Department of Medicine (DAME), University of Udine, Via Palladio 8, 33100 Udine, Italy. <sup>4</sup>The Abdus Salam International Centre for Theoretical Physics (ICTP), Strada Costiera 11, 34151 Trieste, Italy. ✉email: laio@sisssa.it

In this work, we show that a recently developed approach, the Information Imbalance<sup>7</sup>, can be used as a filter to perform feature selection in a clinical framework. We illustrate the procedure on a database of ~ 1300 COVID-19 patients from Udine, including hundreds of features for each patient. These features are extremely heterogeneous, some related with the clinical history, others with the status of the patients at the admission to the hospital, other with the course of the disease, including complications, treatments and clinical outcome. Very importantly, the database is highly incomplete, as is common in clinical databases: the outcome of specific exams (say a TC scan) is typically available only for subsets of the patients, and the clinical history before the admission is often known only partially.

The Information Imbalance approach allows comparing two feature spaces, and deciding if one is more informative than the other. Feature spaces are collections of features that are used to characterize the data. For example, let's say that space A includes age, a specific comorbidity, and the value of a blood test, while space B includes the parameters measured in a TC scan and (also) age. To estimate the Information Imbalance, one finds for each patient their nearest neighbor, which is the other patient that is most similar (closest) according to a distance estimated using the features in space A. In this study we use the Euclidean distance, but the method also allows using other distance measures. Say that for patient number 1 the most similar is patient 412, such that patient 412 has distance rank 0 with respect to patient 1 in space A. Next, one computes the Euclidean distance between patient 1 and 412 in space B, *i.e.* using the features of space B, and finds the number of patients which are closer to patient 1 than patient 412. One repeats this test for all the patients and computes the average of this number. The Information Imbalance, denoted in the following  $\Delta(A \rightarrow B)$ , is proportional to this average. If  $\Delta(A \rightarrow B)$  is small, space A is *predictive* of space B, as patients which are close in A are also close in B, and therefore the average will be taken over small distance ranks. If instead this number is large, the nearest neighbor patients in A are typically "far" in B, which implies that space A is *not informative* of space B.

A key property of this statistic is its asymmetry. At variance with correlation coefficients, the Information Imbalance can have different values when calculated from space A as reference to space B as target and vice versa. This allows capturing standard correlations like measured by Pearson and Spearman correlation coefficients, as well as complex, asymmetric correlations, which are identified by an asymmetric Information Imbalance. In this work we show that Information Imbalance can be adapted to work as a filter to perform feature selection in clinical databases. An important advantage over other filter methods is that most statistical filter methods treat only one variable at a time. They can therefore select subsets of features which are important but redundant<sup>8</sup>. More complex filter methods are designed to integrate variable relevance while minimizing variable redundancy<sup>9</sup>. The approach described in this work automatically selects features which are not only relevant, but also uncorrelated. Indeed, the Information Imbalance can directly be computed for distances including arbitrarily many features. This is an important practical advantage with respect to other methods which are based on comparison between two variables at a time. It also allows comparing the predictive power of subsets of features of different sizes. This, as we will see, allows finding maximally informative subsets of features along with the optimal dimension.

## Results

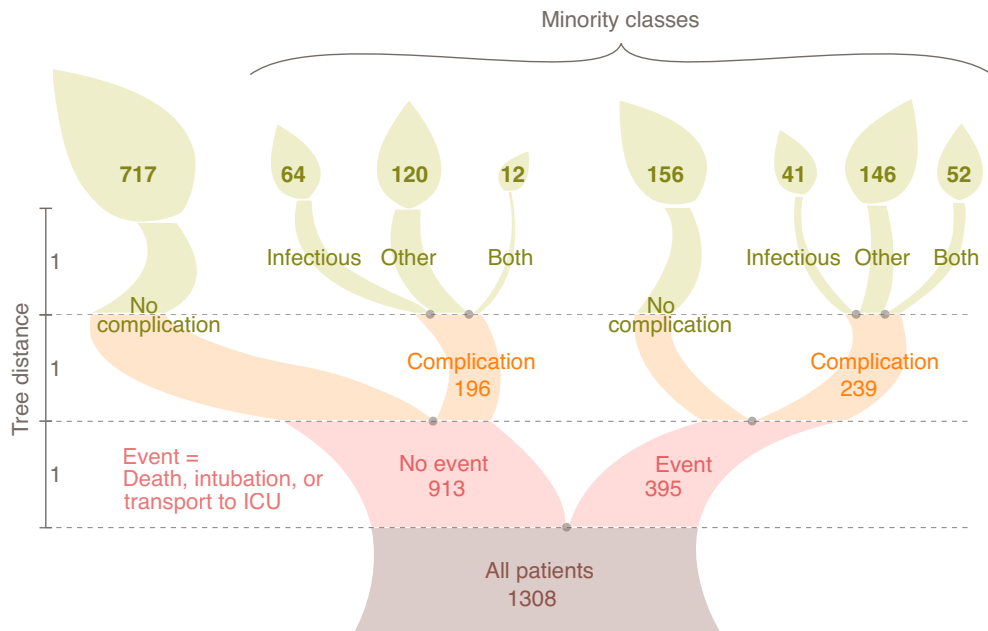
This study is based on a dataset of 1308 COVID-19 patients with ~ 150 features, and 33% of missing values. Some of these features are *input* features, measurable upon admission in the hospital. These include age, gender, physical exams (*e.g.* blood pressure, temperature), blood tests (*e.g.* biomarkers like interleukins), arterial blood gas (*e.g.* partial pressures of oxygen and CO<sub>2</sub>, pH of the blood), chronic comorbidities (*e.g.* diabetes) and chronic medications (*e.g.* diuretics, steroids). 36% of all input features were missing, with only 28 features being complete, and 25 features available in a quarter of the patients or less. The output features were 14, all binary. These features could only be measured later on during the COVID-19 infection, and are therefore the variables that a clinician might be willing to predict. These are available for all patients and include death, intubation, transfer to ICU, and 11 complications (heart attack, pulmonary embolism, arrhythmia, atrial fibrillation, stroke, thrombosis, pneumothorax, pneumomediastinum, hemorrhage, delirium, and secondary infections during hospitalization).

### Feature selection by optimization of the Information Imbalance

We first consider a classical problem in feature selection, finding a small subset of the 138 input features which is maximally predictive with respect to the output features. In the case of the database analyzed in this work, the output features are 14. These features are all binary ("yes" or "no") and are quite heterogeneous in nature. In accordance with the medical insight of the clinicians co-authoring this study, we organized the output features in a "severity tree" (Fig. 1 and Methods).

In short, death, intubation and transfer to ICU were used to split patients into two classes, one for which at least one of such events has occurred (the patients whose course has been more severe), the other in which no event has occurred. The other output features are associated to infectious and non-infectious complications, which are important to decide a clinical strategy. We used the features associated to complications to split the two main classes into four subclasses.

This leads to a classification of patients into eight severity classes. The distance between two patients is then estimated by counting the number of links separating their leaves in the severity tree, divided by two (one directional). For example, the distance between a patient in class one (717 patients, no-event-no-complication) and a patient in class three (120 patients without event but with other complications) is two, while the distance between a patient in class one and in class five (156 patients with event but without complication) is three. This tree distance is the target of the feature selection, however the nominal value of the distances is unimportant and only their relative order matters, since the Information Imbalance method uses distance ranks, *i.e.* the closest neighbor in distance is assigned rank 0, the second closest rank 1, *etc.* The feature selection algorithm presented in this work works as follows: We try to identify a distance A, built as the Euclidean distance using a



**Figure 1.** The severity tree splits patients into eight severity classes of different class size. These classes are depicted here as the leaves. The gray numbers on the left indicate tree distances between nodes. Distances between patients on the tree are measured one-way, such that the distance between patient A in the first big leaf (no event, no complication) to patient B in the third class (no event, other complication, 120 patients) is 2. This number is found by starting from patient A and reaching patient B: We start in leaf 1, take two steps in tree distance down to the level between orange and red, and then take two steps up to patient B's class. Counting this trajectory only downwards, the distance between the patients is 2.

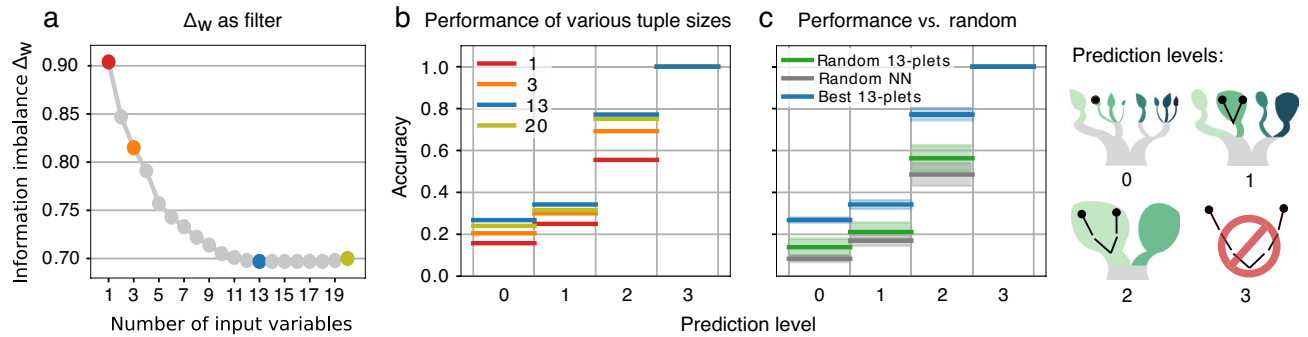
combination of several input features, whose distance ranks are maximally informative with respect to distance ranks  $B$  measured on the severity tree. Degeneracies in input features were treated by addition of small random numbers (Methods). The Information Imbalance between  $A$  and  $B$  is used as a feature selection filter to discriminate between different choices of  $A$  (namely of input features) and select the best one. The classes, “leaves” of the severity tree, are not populated uniformly: class 1 has 717 patients, and the smallest class has only 12 patients. A direct application of classic Information Imbalance in class imbalanced data sets biases class prediction heavily towards the majority class and leads to low class predictive accuracy in the minority classes, like many naive variable selection schemes in classification tasks<sup>10</sup>. Therefore, Information Imbalance<sup>7</sup> has been modified by introducing class weights, aimed at compensating the occurrences of the different severity classes in the data set (Methods). We denote this modified, Weighted Information Imbalance by  $\Delta_w$ .

To identify the best combination of input features, we find the combination of  $n$  input variables minimizing  $\Delta_w$ , which are present in at least 100 patients (Methods). For small  $n$ , the search can be performed exhaustively by testing all the possible combinations of variables. For large  $n$ , the number of possible combinations grows factorially. We use the deterministic beam search algorithm (Methods) which allows finding the best combination of variables for arbitrary  $n$  with great confidence. In Fig. 2a, we plot the optimal value of  $\Delta_w$  as a function of  $n$ . This value decreases up to  $n \approx 13$ , then starts growing slowly, indicating that adding more variables *reduces* the information. This can happen when the new variables only add noise, and no independent information. This analysis indicates that the most informative combination of features includes 13 variables (blue dot in Fig. 2a). The  $n$ -plets of features of different sizes in Fig. 2a are listed in Suppl. Inf. Section 1.

The feature space optimizing  $\Delta_w$  is spanned by these features: *Pathologies*: Liver disease (hepatopathy); *Chronic therapies*: Steroid therapy, potassium sparing diuretics; *Blood exams*: Alanin aminotransferase (GPT), antithrombin III (AT3), interleukin 10 (IL-10), troponin (TROP), antinuclear antibodies (ANA), antineutrophil cytoplasmic antibodies (ANCA), anti smooth muscle antibodies (ASMA), Lymphocytes (FLLINF), percentage of eosinophils (FLEOS%), percentage of neutrophils (FLNEU%).

This combination of 13 features suggests on the one hand a systemic inflammation and autoimmunity, signaled by neutrophils and autoantibodies, and on the other hand an immune paralysis and anti-inflammatory effort (*i.e.* steroid therapy, IL-10). Furthermore, it has already been suggested that the up-regulation of inflammatory markers can lead to the progression of the disease to the severe form and eventually cause liver damage in these patients<sup>11</sup>, which suggests why hepatopathy might be an important feature.

The information provided by these input features on the severity of the course of the disease is assessed by using these variables to predict the class of each patient. We first use a prior-corrected  $k$ -NN classifier, in which the class of a patient is assumed to be the same of their 10 nearest neighbors according to the input features (see Methods for details). This predictor has no variational parameters, and lacks therefore tunability, but allows assessing directly the consistency within the neighborhood of the patients induced by our optimization



**Figure 2.** (a) The optimal (lowest) Information Imbalances  $\Delta_w$  as a function of the number of input features. (b) Accuracies of 10-NN predictions at given prediction levels for tuple sizes as marked in (a) (the line shows the average over the 10 best  $\Delta_w$  results for each tuple size, see main text). The accuracy corresponds to the fraction of patients predicted correctly at a given prediction level. The prediction level corresponds to the maximum tolerated distance of the true vs. the predicted class on the severity tree, as depicted on the right. Together, the lines of one color can be considered the CDF of the fraction of patients predicted correctly. (c) Accuracies of the top 13-plets, benchmarked vs. randomly drawn 13-plets (green), and random nearest neighbor assignments (grey). The average over the 10  $\Delta_w$  results is shown as bold line and the standard deviation as shade. The predictions were generated by LOO. The graphic on the right explains prediction levels: 0 denotes the fraction of patients for whom the correct class was predicted; 1 indicates the fraction for whom the occurrence of event and complication (yes or no) was predicted correctly; 2 means the patients predicted on the correct tree side (event vs. no event). The remaining fraction of patients was predicted wrongly (distance = 3).

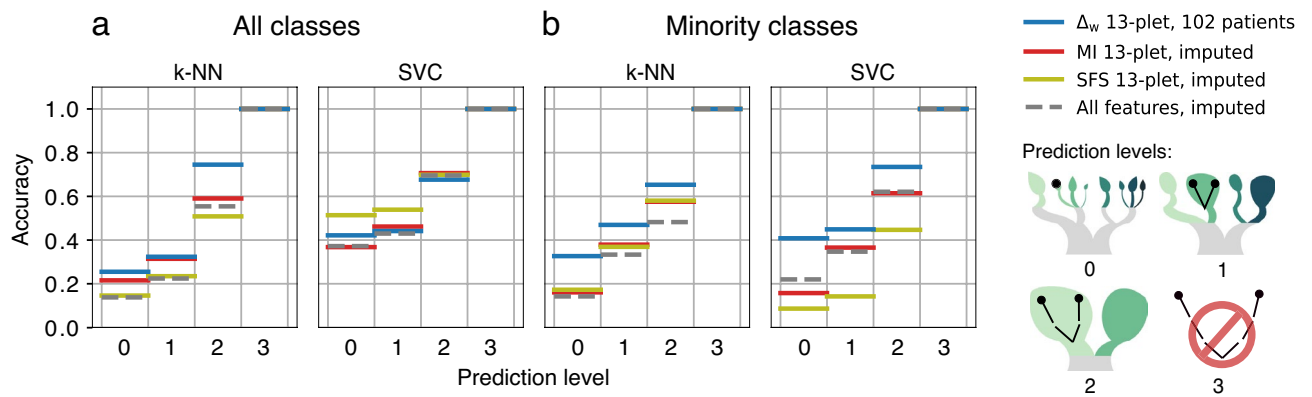
procedure. Later in this study, we also use a support vector classifier (SVC) to compare the results to the k-NN classification. In Fig. 2b, we plot the accuracies of the prediction at different prediction levels, and for a different number of features ranging from 1 to 20. The accuracy corresponds to the fraction of patients predicted correctly at a given prediction level, defined by the maximum tolerated distance of the true vs. the predicted class on the severity tree. Level 0 is the fraction of patients predicted completely correctly, level 1 is the correct prediction of having an event and complication, and level 2 is the fraction for whom only the event was predicted correctly (correct tree side). The remaining percentage belongs to patients misclassified according to having an event or not (level 3). The accuracies are estimated by a leave-one-out (LOO) validation procedure. The performance increases with the number of features, as seen in the height of the CDF curves. This effect staggers with bigger tuple sizes and levels out.

Our approach automatically also leads to a selection of features which are practically uncorrelated with very small pairwise Pearson correlation coefficients (Supp. Inf. Section 5.1). The mean of the pairwise correlations of numerical features in the best 13-plet is  $\bar{r} = 0.02$ , and the pairwise Information Imbalances mean is  $\bar{\Delta} = 0.96$ . This happens “for free” since adding a feature which can be predicted by other already selected features does not significantly improve  $\Delta_w$  (see Supp. Inf. Figure S4).

The Weighted Information Imbalance of the top-ranking  $\Delta_w$ -optimized tuples of the same size is very similar. E.g. the best 13-plet has  $\Delta_w = 0.69$ , while the tenth best 13-plet has  $\Delta_w = 0.70$  and differs by only 3 features from the best. Therefore, in Fig. 2b and c the plotted lines are averages over the top ten results.

To put these results in context, we need to compare them to a baseline. We first performed a comparison with the predictive performance of randomly selected tuples. Secondly, we performed a comparison with a prediction performed by assigning each patient a “nearest neighbor” at random. To be comparable to our  $\Delta_w$ -optimized tuples, both comparisons use averages over the ten best performing random 13-plets (Fig. 2c with shaded standard deviation). The prediction accuracy is much higher in the  $\Delta_w$ -optimized 13-plets than when using random variables. With the  $\Delta_w$ -optimized 13-plets the exact target class predictions are about 27% of cases (distance = 0), while predicting the correct side of the severity tree (event vs. no event, distance = 2) has a quota of around 77%. In comparison, the completely random result distinguishes event vs. no event with under 50% accuracy and the random 13-plets from our data set predict the same with an accuracy of under 60%.

We then compared our results to standard feature selection methods, namely to selection by a mutual information (MI) score and to forward sequential feature selection (SFS). MI is a filter method which ranks the individual features against the output classes, while SFS uses the predictor method (here: k-NN and SVC) in a greedy approach to search the best n-plets, and is hence a wrapper method (see Methods). Using the top ranking 13 features selected by MI, the 13-plet selected by SFS (different for k-NN and SVC), the 13-plet selected by our approach, and all features, we perform a prediction with these four sets of features using two different approaches, the k-NN predictor with k=10 and a SVC predictor with balanced class weights. Unlike Information Imbalance, these two feature selection methods do not provide a way to select the optimal tuple size at the feature selection stage. Therefore, the predictive performances of optimal tuples of several sizes selected with MI and SFS was compared (Supp. Inf. Figure S3). Based on these results, the 13-plets were chosen, since in all three feature selection models this size is optimal or nearly optimal. The results are presented in Fig. 3a. Using the k-NN predictor the accuracy of the prediction is significantly higher if one uses our approach. Using SVC, SFS performs better than the other feature selection methods, especially at a prediction level 0 (Fig. 3a SVC). However, doing the prediction only for the seven minority classes, i.e. excluding the no-event-no-complication



**Figure 3.** Accuracy of prediction using the top  $\Delta_w$ -optimized 13-plet, vs. using the 13 features with the highest mutual information (MI, filter) score on the imputed data set (red) and the 13-plet selected by forward sequential feature selection (SFS, wrapper) on the imputed data set (olive). The dashed grey line shows the prediction using the complete, imputed dataframe without prior feature selection. LOO was used for generating the predictions. Accuracy corresponds to the fraction of patients predicted correctly at a given prediction level. The prediction level corresponds to the maximum tolerated distance of the true vs. the predicted class on the severity tree, as depicted on the right. **(a)** Accuracies, using k-NN (10-NN) and SVC predictors, respectively. **(b)** The same as **(a)**, but considering the predictions of the seven minority classes only (excluding no-event-no-complication). The SFS 13-plets for k-NN and SVC prediction are two distinct ones, using k-NN and SVC predictors accordingly in the SFS-construction of the 13-plets.

class (Fig. 3b), the feature tuple obtained with our approach consistently outperforms all other tuples in both, k-NN and SVC predictions. The SFS tuple has especially low accuracy for minority classes in the SVC prediction, even though balancing class weights were employed in the tuple generation and prediction. We recall that in imbalanced multiclass prediction the problem is twofold: prediction in imbalanced datasets tends to favor the majority class, and on top of this, prediction is more complex than in binary prediction, because there can be several majority and minority classes with various relationships towards each other<sup>12,13</sup>. Furthermore, the error introduced in the imputation could effect the minority classes more, as previously elaborated for standard imputation methods<sup>14</sup>.

Their reliance on imputation is an Achilles' heel of both, NMI and SFS, because most standard implementations, such as scikit-learn<sup>15</sup>, need data sets without missing values. In our dataset, slightly more minority class data had to be imputed than majority class data (38% vs. 35%) and the different feature selection methods selected tuples with 49% (NMI), 19% (SFS with 10-NN) and 32% (SFS with SVC) imputed data. Weighted Information Imbalance, as introduced in this work, uses a constrained, data-set-reductive approach (see Methods) and hence does not impute features. It is a non-parametric algorithm employing class-balancing weights which can intrinsically handle multiclass groundtruths.

The prediction accuracies were also compared with two regularized classifiers,  $L_1$  (lasso) regularized logistic regression and regularized sparse SVC (supplimentary Fig. S2). In terms of general predictivity of all classes these methods do well: sparse SVC even reaches over 50% at the class level. This effect, however, does not translate to the minority classes, where this same classifier only reaches 19% correct predictions, while Information Imbalance selection and subsequent SVC reach 41%.

All in all, given the importance of minority classes in the clinical setting, we conclude that the Information Imbalance filter method finds a superior feature subspace for severity prediction in the present COVID-19 database than the two other feature selection methods considered here, and also then state-of-the art regularized classifiers, considering the performance in minority classes.

### Identifying important but rarely available features

The 13 features which have been rarely identified as optimal as described above are simultaneously available for only 102 patients. Indeed, approximately one third of the data are missing in the data set. Moreover, missing values are not evenly distributed among the features: some “cheap” exams are performed routinely for all the patients, others are performed only for a small fraction of the patients. As a consequence, optimal features might not be available for a generic patient.

We repeated the optimization procedure by attempting to select, independently for each patient, the best tuple among the features which are available for them. We beam-searched over the possible feature tuples for the respective patient, using the subset of the patients which had no missing values in these features (see Supp. Inf. Section 2). The patient-specific optimal tuple was subsequently used in a k-NN prediction, using the other patients with full information as training set. The average predictive performance of the features selected in this manner is still very significant but is reduced as compared to what is observed for the optimal 13-plets (Supp. Inf. Figure S1). The prediction of the correct side of the tree (event vs. no event) is reduced by about 7% to roughly 70% and the correct class prediction drops from 27% to about 23%. This result is not surprising: the variables which turn out to be most informative happen to be simultaneously available only for a relatively small fraction of patients.

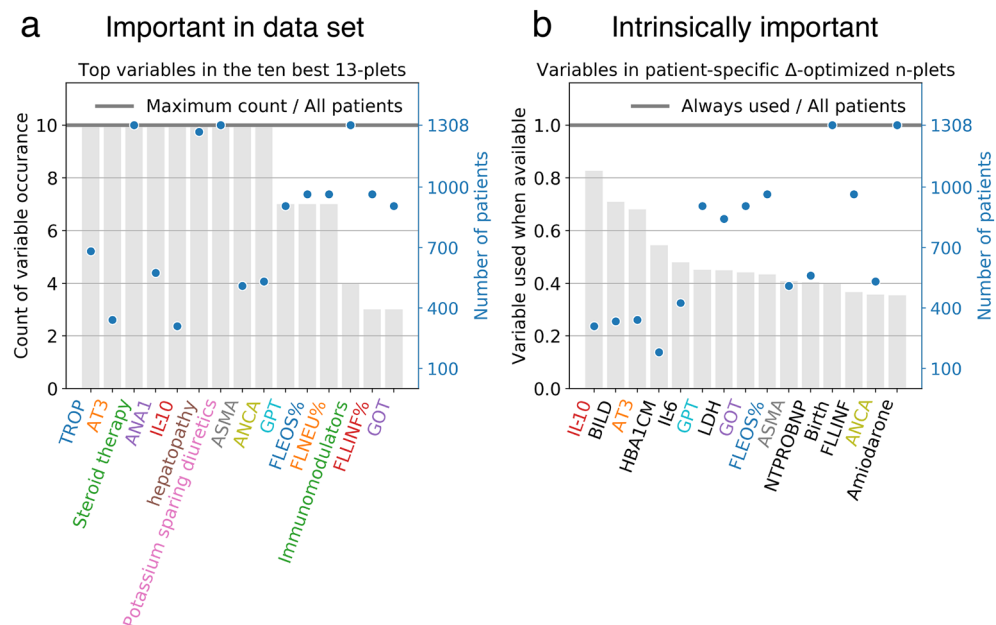


This result pushed us to develop a quality measure for the input features which takes into account the fact that  $\Delta_w$ -optimized n-plets contain features that are “good” in two ways: firstly, they are intrinsically important, and secondly, available together in the same patients. We estimate, for each feature  $f$ , the number of patients  $NP_f$  for which  $f$  is included in the most informative  $\Delta_w$ -optimized tuple (which is patient-specific). The usage  $U_f$  of feature  $f$  is then estimated by the ratio between  $NP_f$  and the number of patients for which  $f$  is available. If  $U_f$  is close to one, then the feature has been chosen for each patient where it was available, and has a high intrinsic importance for the Information Imbalance towards the severity tree target classes.

Figure 4a shows the features which are used in the 10 most predictive 13-plets (predictive performance of these in Fig. 2c) which, we recall, are present at the same time in  $\sim 100$  patients. 9 of the 13 features are used in all the ten best models. The blue dots indicate the fraction of patients for which the variable is available. For example, information about steroid therapy, which is used in all the 10 best models, is available for all the patients. AT3, also used in all the best models, is available for approximately 350 patients. Figure 4b on the other hand shows the value of the “usage when available” statistic  $U_f$ . Some of those variables are present in both sets of Fig. 4, namely the cytokine IL-10 (interleukin 10), the anticoagulant protein AT3 (antithrombin III), the autoantibodies ANCA and ASMA (antineutrophil cytoplasmic antibody and anti-smooth muscle antibody), the liver enzymes GPT and GOT (alanin aminotransferase and aspartate aminotransferase), and the percentage of eosinophils, FLEOS%. The  $U_f$  statistic can also find intrinsically good predictors which are underrepresented due to missing values, and as such do not appear in the most used features of the  $\Delta_w$ -optimized tuples. For the COVID-19 severity prediction some of these are direct bilirubin (BILD), the diabetes indicator glycated hemoglobin (HBA1CM), interleukin 6 (IL-6), and the enzyme lactate dehydrogenase (LDH), which indeed have a high value of  $U_f$  (panel b) but do not enter in the best model (panel a). IL-6 is a pro-inflammatory cytokine and has previously been linked with COVID-19 severity<sup>16</sup>. Abnormal bilirubin levels indicate sepsis, and the severity of patients could be linked to the fact that they have developed sepsis, and therefore a condition of a dysregulated systemic response. Glycated hemoglobin is linked to a condition of decompensated diabetes, which predisposes to infections. Diabetes is known to predispose severe COVID-19 infections<sup>17</sup>.

## Discussion

In this work, we show that the Information Imbalance<sup>7</sup> approach can be used to perform feature selection in a clinical database. We illustrated our approach by the analysis of a database of 1300 COVID-19 patients from Udine hospital with  $\sim 150$  features for each patient and one third of missing data. In order to deal with unbalanced classes in a clinical setting, we had to include weights in the Information Imbalance definition. We find that the optimal feature tuples selected by our approach perform better in the k-NN prediction of COVID severity classes than two other standard feature selection methods, mutual information (a filter) and sequential feature selection (a wrapper), implemented in a standard statistical analysis package<sup>15</sup> (see Fig. 3). Regulated classifiers which do not include a prior feature selection step can reach slightly higher accuracies for the prediction of all classes (Supp. Inf. Figure S2). This effect, however, does not translate to minority classes, which in our database



**Figure 4.** (a) Statistics of the features present in the best 10 13-plets. The grey bar indicates the fraction of 13-plets in which the variable is present. The blue dots indicate in how many patients the feature is available (scale defined on the right y-axis). (b) Intrinsically important features estimated by the “usage when available” statistic  $U_f$ . The gray bars indicates the value of  $U_f$ , which is large when a variable is always used when present. The blue dots are the same as in panel A.

correspond to patients developing very severe symptoms. Information Imbalance feature selection followed by SVC classification outperforms all other methods in identifying those patients.

The classification task, as applied in our study, serves to validate the utility of the  $\Delta_w$  selected features in distinguishing between severity classes of COVID-19, rather than introducing a complete novel classification model. A future direction of research could extend this work into a framework including techniques for quantifying the uncertainty of prediction, but is beyond the scope of the current work and would overshadow the core contribution of our novel feature selection method.

The optimal variables of our data set include cytokines (such as IL-10), autoantibodies (ANA, ASMA, ANCA), and therapies that reduce the immune response (steroid therapy and immunomodulators as chronic home therapy). These findings are of medical interest because the pathogenic mechanisms that drive COVID-19 clinical deterioration can likely be contributed to systemic inflammation, disordered coagulation (AT3), and immune dysfunction. The cytokine storm that characterizes the unfavorable outcome of patients comprises classical markers of systemic inflammation such as IL-6, which is now largely disposable as a single diagnostic test also at urgent request in the majority of hospitals around the world. IL-6 increases during COVID-19 illness decline as patients recover, correlating with the severity of the disease course<sup>16</sup>. When IL-6 levels are already very high, the focus can be shifted to the degree of immunoparalysis and anti-inflammatory effort (IL-10)<sup>16</sup>. COVID-19 is known to alter the coagulation state and, in severe cases, lead to hypercoagulation, which is causally involved in negative patient outcomes. AT3 levels decrease in inflammatory conditions and AT seems to possess anti-viral properties<sup>18</sup>. The role of transaminases (GPT and GOT), and history of liver disease are also interesting. The hepatic consequences of an SARS-CoV-2 infection are recognized as an important component of COVID-19 and this aspect is most clinically relevant in patients with pre-existing cirrhosis<sup>19</sup>. There are several other potential contributors to abnormal liver biochemistries in COVID-19, including ischaemic hepatitis, hepatic congestion related to cardiomyopathy, and transaminase release due to the breakdown of skeletal and cardiac muscle<sup>20</sup>.

As common in most real world clinical data, also in this case the accuracy does not allow to perform exact prediction of the patient fate into their severity class (prediction level 0 in Figs. 2 and 3). Exact predictions, as tested in a leave-one-out-cross-validation approach, happen in 30% - 40% of cases, depending on which predictor is used and which classes is considered. However, in over 70% of the cases we were able to predict if the patient will suffer a serious event (death, transport to ICU or intubation, prediction level 2) or not. In this sense, the hierarchical tree structure of the output space is of advantage, because even if the exact prediction is not possible, a warning along a less stratified prediction level of the output tree is nevertheless possible. The feature tuples derived from our approach keep the high accuracy described above also for minority class patients, as opposed to tuples generated with other feature selection techniques.

One important advantage of the approach presented here is that it works without imputation<sup>21</sup>, namely it does not require the preprocessing step of assigning missing values of the input features. While MI and SFS in their standard implementations need complete data sets and such require imputation, our approach finds patient-specific optimal input feature tuples in the original, incomplete data.

A limitation of the method is presented by the nature of the ground truth space. Since Information Imbalance finds input feature spaces which reproduce the neighborhood relationships observed in a ground truth space, it works best if the ground truth is either continuous, or at least has classes for which one can meaningfully identify a distance similar to our severity tree. The method is less suitable for learning a ground truth distance which can take only a few values, and would not be appropriate for binary classification tasks. Furthermore, continuous input features are more likely to be chosen in small tuples, because they can carry more information than categorical variables. The optimal 13-plet of this study includes three binary variables (hepatopathy, steroid therapy, potassium sparing diuretics: yes/no), proving that there are several informative binary features in this COVID-database that hold information complementary to the chosen numerical features.

The threshold of  $\geq 100$  patients without missing data and Jensen-Shannon divergence of  $\leq 0.06$  for Information Imbalance feature selection were a practical choices, which can influences the specific results of the optimal tuple. While this choice has proven to balance performance and representativeness effectively, different informative feature tuples could be selected for different thresholds. The optimal 13-plet should be considered one out of several possible informative feature combinations in this data set.

This work also tested the features for their intrinsic importance decoupled from their availability in the data set by employing a simple usage statistic, which we called  $U_f$ . This analysis can be used to provide recommendations to clinicians for future data collections, because it identifies potentially important features for COVID severity prediction, which are, however, not abundant enough in our data set. For example, we find that conjugated (direct) bilirubin (BILD), is available for less than 30% of the patients in the current data set, yet it is selected very often, when present, for the patient-optimal prediction. A similar scenario is found for glycated hemoglobin (HBA1CM) and interleukin 6 (IL-6). The collection of these features should be emphasized in future data collection efforts.

The use of clinical features and diagnostic features as biomarkers is of great clinical interest, in order to facilitate improved triaging and earlier therapeutic decisions. The model presented here could help the clinicians to focus on the variables of greatest interest in order to target the allocation of resources and escalation of care.

## Methods

### The data set

This work is a retrospective study involving data of 1308 COVID-19 patients from Udine hospital. The data set includes patients admitted to the Infectious Disease ward of the Azienda Sanitaria Universitaria Friuli Centrale Santa Maria della Misericordia of Udine, a 1000-bed tertiary-care teaching hospital identified as a regional

referral center for COVID-19 patients, from March 2020 to March 2021. Informed consent was obtained from all participants (see section Informed consent).

For all patients the following parameters were collected: evaluation of in/exclusion criteria; socio-demographics (age, gender, race, height, weight); date and time of the onset of symptoms and of the admission to the hospital; ward of hospitalization; co-morbidities (dyslipidemia, obesity, diabetes, chronic obstructive pulmonary disease, chronic kidney injury, liver disease, hypertension, solid and hematologic neoplasms, autoimmune diseases, primary or secondary immunosuppression), including Charlson score index; findings from routine physical examination (temperature, heart rate, breathing rate, blood pressure, SPO<sub>2</sub>, neurological status); routine diagnostics performed (chest X-ray, CT scan, ultrasound, microbiological tests and blood tests performed); date and time of blood sampling initial and final diagnosis; type and focus of infection; date of discharge; date and time of ICU admission and discharge; needs for organ support and/or invasive ventilation; any serious adverse event or complication which occurred during hospitalization; therapies carried out; lab parameters from routine blood testing which were assessed at presentation (within 48 hours of admission); data from blood gas analysis, such as PaO<sub>2</sub>/FiO<sub>2</sub> ratio, alveolar arterial gradient, and lactate.

The features were divided into input (measurable upon hospitalization) and output (severity of outcome of COVID infection) features. 14 output features were decided upon by medical knowledge. The natural hierarchy in the severity of these features led to the creation of a “severity tree”, where death, intubation and transfer to ICU were used to split patients into an “event” group, where at least one of these three events occurred, and a “no event” group. Various infectious and non-infectious complications make up the second split of the tree, as seen in Fig. 1. These output features were available for all patients, yielding a distance (and an identical distance rank) between 0 and 3 for each pair of patients on the tree. Since there are only eight classes in the severity tree, the distances between all patients in this output space are degenerate, meaning that many patients have the same distances from each other. Furthermore, there is a high class imbalance, *i.e.* the biggest class contains more than half of the patients.

138 input variables were selected, including age, gender, physical exams (*e.g.* blood pressure, temperature), blood test (*e.g.* diagnostic antibodies and interleukins), hemogas (*e.g.* partial pressures of oxygen and CO<sub>2</sub>, pH of the blood) values and chronic comorbidities (*e.g.* diabetes) and medications (*e.g.* diuretics, steroids). Potential input spaces are tuples of these input variables. As many real world data, the input data is characterized by missing values, which leads to reduced patient numbers for certain combinations of input features. In some possible input spaces also distances between data points are degenerate, due to categorical, binary and repeated values.

### The Information Imbalance

The Information Imbalance  $\Delta$  is defined as follows (for more details see<sup>7</sup>):

$$\Delta(A \rightarrow B) \approx \frac{2}{N} \langle r^B | r^A = 1 \rangle. \quad (1)$$

This means that  $\Delta$  is a scaled mean of the distance rank of a specific neighbor of each data point in space  $B$  ( $r^B$ ), where that specific neighbor is the nearest neighbor (has distance rank 1,  $r^A = 1$ ) of the data point in space  $A$ . Here, a data point is a patient. The rank is the distance position of one point to another in relation to all other points, *i.e.* the nearest neighbor of a point has rank 1 to that point, the second nearest neighbor rank 2, in the original formulation. Please note that due to the tree structure in the COVID-19 output space in this paper we assign the nearest neighbor rank 0, which leads to very small numeric differences in the case of few data points. Several distance metrics could be used to calculate these ranks. Since the method is focused on the identification of a feature space which reproduces the nearest neighbors of another feature space, it is not very sensitive to the precise choice of the distance metric. While the distance between two ‘far’ points will likely be very different if computed, *e.g.*, with the Hamming distance or with the Euclidean metric, the nearest neighbors are more preserved across metrics. This work uses the Euclidean distance.  $\langle \cdot \rangle$  denotes the expectation value, the arithmetic mean.  $N$  is the number of all data points, in this case all patients. Equation 1 is estimated as

$$\Delta(A \rightarrow B) = \frac{2}{N} \frac{\sum_{i=1}^N r_i^B}{N}, \quad (2)$$

where  $r_i^B = r_i^B$  given  $r_i^A = 1$ . The Information Imbalance between feature space  $A$  and feature space  $B$ ,  $\Delta(A \rightarrow B)$ , is proportional to the average of the neighbor ranks in space  $B$ , conditioned to nearest neighbors in space  $A$ , and normalized such that if  $A$  predicts space  $B$  perfectly,  $\Delta(A \rightarrow B) \approx 0$ , and if  $A$  has no information on  $B$ ,  $\Delta(A \rightarrow B) \approx 1$ .

### Class-corrected weighted Information Imbalance

For the high dimensional classification of COVID-19 severity, Information Imbalance is used in the feature selection step. The task is to find input variable spaces in which the nearest neighbors optimally describe the distribution of the output classes. The severity tree output space is degenerate with high class imbalance, where leaf 1 has more than half of the total patients. Using naive classifiers in class imbalanced data sets biases class prediction heavily towards the majority class, especially when feature selection is performed. Subsequently, the class predictive accuracy is low in the minority classes<sup>10</sup>. The original implementation of Information Imbalance was developed for continuous in- and output spaces and has the described shortcomings when applied to few, imbalanced output classes. However, in sick patient prediction there is a high cost associated to false negatives. To find feature spaces which also predict the neighborhood of the patients in small classes, we introduce class weights  $w_{i,l} = \frac{1}{\text{leafsize}_l}$ , for each patient  $i$  in class  $l$ . Furthermore, the distances in the severity tree are highly



degenerate. The structure of our tree leads to a total of four distinct distances and identical four distance ranks, where rank 0 means that two patients are in the same class and rank 3 means the two patients are on opposite sides of the tree (event vs. no event). Since the classes are imbalanced, the average probability to find a certain rank neighbor from the different classes is not uniform. Therefore, the normalization  $a$  is built to reflect this and bring the average value of  $\Delta$  to a value of 1, when the nearest neighbor ranks are distributed randomly. The adjusted “Weighted Information Imbalance  $\Delta_w$ ” becomes:

$$\Delta_w(A \rightarrow B) \approx a \frac{\sum_{i=1}^N r_i^{B'} w_i}{\sum_{i=1}^N w_i} \quad (3)$$

If more than one nearest neighbor exists at the same distance in the output space B, the nearest neighbor rank is the mean over these  $M$  nearest neighbors of patient  $i$ :  $r_i^{B'} = \frac{\sum_{j=1}^M r_{ij}}{M}$ .

For this implementation the input space needs to provide clear nearest neighbor assignments, in order to find the according ranks in the output tree. To resolve the degeneracy in the input space, small random numbers are added to duplicated input values. Since this makes the estimated imbalance a stochastic variable, we repeated the optimization ten times with different random seeds, verifying that the results are robust. For the ten implementations, the Weighted Information Imbalances of optimal tuples with the same tuple sizes are mostly identical, up to the second digit (equal to Fig. 2a) with standard deviations on the order of  $10^{-4}$ . Also the chosen tuples for each size are largely congruent, with the best single variable always being brain natriuretic peptide (BNP) and the best 13-plets of the ten implementations being identical to the one present in this study, except in one case, where eosinophils (FLEOS) have been selected instead of lymphocytes (FLLINF).

Variables were normalized by dividing them by their standard deviation in order to move them into a comparable value range.

The problem of missing data was treated in a constrained, data-set-reductive manner, by which only feature tuples were considered which were present in at least 100 patients. Then  $\Delta_w$  was calculated for the feature tuple in question using only these  $\geq 100$  data points with the additional constraint that the base-2 Jensen-Shannon divergence of the tuple class distribution towards the full class distribution (of the 1308 patients) was  $\leq 0.06$ , in order to ensure proportionate stratified sampling, *i.e.* such that the share of the classes in the sub-sample is proportionate to the full sample. This subset of  $\geq 100$  patients has no missing values and can hence be passed on to the downstream classifiers, kNN and SVC, without problems.

Input spaces with more than 1 variable were selected by applying  $\Delta_w$  to a pool of candidate feature tuples selected by beam search with beam width 55. Beam search is a heuristic algorithm<sup>22</sup> which is employed because the full exhaustive search of all possible variable  $k$ -tuples (out of  $D$  features) would lead to combinatorial explosion with complexity  $\mathcal{O}(D^k)$ . Like in the vanilla greedy approach, the pool of candidate feature tuples is sequentially extended by adding new features to the best scoring previous tuples. However, in beam search, not only the one best result is chosen and variables added to it, but the  $n$  best results, where  $n$  denotes the beam width. In this case, the 55 best results were iteratively extended. The computational complexity of beam search is  $\mathcal{O}(D \cdot n \cdot k)$ <sup>23</sup> for beam width  $n$ .

### Prior-corrected k-NN prediction of severity

After finding  $\Delta$ -optimized sets of features, we use an adjusted k-nearest neighbor (k-NN) prediction in a leave-one-out (LOO) approach. Comparing the predicted severity class to their actual class, we evaluate the performance of the method with cumulative distribution functions (CDFs) of the distances  $d$ , *i.e.* consider the fractions of cases in which the class was predicted correctly ( $d = 0$ ), or a neighboring class was predicted ( $d = 1$ ), or the same side of the severity tree was predicted ( $d = 2$ ). The empirical probability distribution of classes for each patient calculated from their nearest neighbors' classes cannot be taken at face value due to the class imbalance. Effectively, the class imbalance makes it much more likely to find a majority class NN than to a minority class NN. The average global density of minority class points is smaller. For this reason we divide the empirical probabilities by the prior probabilities  $P^0$  of the classes in the samples, to create a metric which measures how much bigger or smaller the local density of the various classes is around the patient, in comparison to the average global density. If this value is greater than 1 for a class, this class is more likely than average to be the class of the patient. Since several classes can have values greater than one, the prediction is based on the class with the maximum value.

$$P_{i,l}^0 = \frac{\text{leafsize}_{i,l}}{\sum_{k=1}^N \text{leafsize}_{i,k}} \quad (4)$$

$$\rho_{i,l} = \frac{P_{i,l}^{\text{emp}}}{P_{i,l}^0} \quad (5)$$

### Identifying important but rarely available features by “usage when available”

Due to the missing values in the input data set, the globally selected best variables are a function of their intrinsic goodness of prediction as well as their availability, especially together with other orthogonal features. To decouple these effects to find intrinsically valuable features, the optimal tuple for each patient is found using  $\Delta_w$  in a leave-one-out (LOO) procedure. The “usage when available” statistic  $U_f$  for each feature  $f$  is simply defined as the ratio of the number of times a feature is used in all patient-specific  $\Delta_w$ -optimized tuples  $n_{f,\Delta_w}$ , over the count of the availability of that feature across all patients  $a_f$ :

$$U_f = \frac{n_{f,\Delta_w}}{a_f} \quad (6)$$

### Mutual information and sequential feature selection

Two other standard feature selection methods are compared to Information Imbalance. A frequently used filter is the estimated mutual information (MI) between each feature and the output classification<sup>24</sup>, and a straight forward wrapper method is sequential feature selection (SFS). For both the implementation in scikit-learn<sup>15</sup> is employed and for both methods the data set has to be complete. Thus missing values were filled in by imputation (scikit-learn KNNImputer with 10 NN, uniform weights and Euclidean distance). The MI between each feature and the output classes were calculated with the scikit-learn class `mutual_info_classif` (3-NN). Forward SFS was calculated using the wrapper `SequentialFeatureSelector` (5-fold, 10-NN) and SVC (balanced class weights) respectively). The prior-corrected k-NN predictor was used for classification in a LOO approach as explained above, and compared with a support vector classification (SVC) LOO prediction as implemented in scikit-learn (`sklearn.svm.SVC`), using default settings and balanced class weights.

### Informed consent

From March 2020 to March 2021, clinical data were collected in the retrospective registry “Management coronavirus Disease In hospital registry” - MANDI registry - authorized by the regional ethics committee Comitato Etico Unico Regionale del Friuli Venezia Giulia (CEUR FVG) (decree n. 957, 10/09/2021). Inclusion criteria included age of at least 18 years and given informed consent to use the de-identified clinical data. All research was performed in accordance with the Declaration of Helsinki.

### Data availability

The COVID-19 patient data extracted from the MANDI registry, used for these analyses, as well as the code of the Weighted Information Imbalance optimization, is available at Open Science Framework: [https://osf.io/b4ktd/?view\\_only=2a71edf707bc4df78cb4d3b84c889461](https://osf.io/b4ktd/?view_only=2a71edf707bc4df78cb4d3b84c889461)

Received: 19 July 2023; Accepted: 4 May 2024

Published online: 10 May 2024

### References

- Chandrashekar, G. & Sahin, F. A survey on feature selection methods. *Comput. Electr. Eng.* **40**, 16–28. <https://doi.org/10.1016/j.compeleceng.2013.11.024> (2014).
- Yu, K. *et al.* Causality-based feature selection: Methods and evaluations. *ACM Comput. Surv.* <https://doi.org/10.1145/3409382> (2020).
- Witten, D. M. & Tibshirani, R. Covariance-regularized regression and classification for high dimensional problems. *J. R. Stat. Soc. Ser. B: Stat. Methodol.* **71**, 615–636. <https://doi.org/10.1111/j.1467-9868.2009.00699.x> [https://academic.oup.com/jrsssb/article-pdf/71/3/615/49686350/jrsssb\\_71\\_3\\_615.pdf](https://academic.oup.com/jrsssb/article-pdf/71/3/615/49686350/jrsssb_71_3_615.pdf). (2009).
- Ravikumar, P., Lafferty, J., Liu, H. & Wasserman, L. Sparse additive models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **71**, 1009–1030. <https://doi.org/10.1111/j.1467-9868.2009.00718.x> (2009).
- Xu, S., Bu, Z., Chaudhari, P. & Barnett, I. J. Sparse neural additive model: Interpretable deep learning with feature selection via group sparsity. In: Koutra, D., Plant, C., Gomez Rodriguez, M., Baralis, E. & Bonchi, F. (eds.) *Machine Learning and Knowledge Discovery in Databases: Research Track*, 343–359 (Springer Nature Switzerland, Cham, 2023).
- Guyon, I. & Elisseeff, A. An introduction to variable and feature selection. *J. Mach. Learn. Res.* **3**, 1157–1182 (2003).
- Glielmo, A., Zeni, C., Cheng, B., Csányi, G. & Laio, A. Ranking the information content of distance measures. *PNAS Nexus* <https://doi.org/10.1093/pnasnexus/pgac039> (2022).
- Kuhn, M. & Johnson, K. *Applied Predictive Modeling* (2013).
- Fleuret, F. Fast binary feature selection with conditional mutual information. *J. Mach. Learn. Res.* **5**, 1531–1555 (2004).
- Blagus, R. & Lusa, L. Class prediction for high-dimensional class-imbalanced data. *BMC Bioinform.* <https://doi.org/10.1186/1471-2105-11-523> (2010).
- Amiri-Dashatan, N., Koushki, M., Ghorbani, F., Naderi, N. Increased inflammatory markers correlate with liver damage and predict severe COVID-19: a systematic review and meta-analysis. *Gastroenterol. Hepatol. Bed Bench* **13**, 282–291. <https://doi.org/10.22037/ghfbb.v13i4.2038> (2020).
- Tanha, J., Abdi, Y., Samadi, N., Razzaghi, N. & Asadpour, M. Boosting methods for multi-class imbalanced data classification: an experimental review. *J. Big Data* **7**, 1–47 (2020).
- Li, Q., Song, Y., Zhang, J. & Sheng, V. S. Multiclass imbalanced learning with one-versus-one decomposition and spectral clustering. *Expert Syst. Appl.* **147**, 113152. <https://doi.org/10.1016/j.eswa.2019.113152> (2020).
- Awan, S. E., Bennamoun, M., Sohel, F., Sanfilippo, F. & Dwivedi, G. Imputation of missing data with class imbalance using conditional generative adversarial networks. *Neurocomputing* **453**, 164–171. <https://doi.org/10.1016/j.neucom.2021.04.010> (2021).
- Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830. <https://doi.org/10.48550/arXiv.1201.0490> (2011).
- Fabris, M. *et al.* Cytokines from bench to bedside: A retrospective study identifies a definite panel of biomarkers to early assess the risk of negative outcome in COVID-19 patients. *Int. J. Mol. Sci.* <https://doi.org/10.3390/ijms23094830> (2022).
- National Center for Immunization and Respiratory Diseases (NCIRD), Division of Viral Diseases. People with certain medical conditions. <https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/people-with-medical-conditions.html>. Accessed: 2022-03-18.
- Schlömmer, C., Brandtner, A. & Bachler, M. Antithrombin and its role in host defense and inflammation. *Int. J. Mol. Sci.* **22**, 4283. <https://doi.org/10.3390/ijms22084283> (2021).
- Marjot, T. *et al.* COVID-19 and liver disease: Mechanistic and clinical perspectives. *Nat. Rev. Gastroenterol. Hepatol.* **18**, 348–364. <https://doi.org/10.1038/s41575-021-00426-4> (2021).
- Jothimani, D., Venugopal, R., Abedin, M. F., Kaliamoorthy, I. & Rela, M. COVID-19 and the liver. *J. Hepatol.* **73**, 1231–1240. <https://doi.org/10.1016/j.jhep.2020.06.006> (2020).

21. Donders, A. R. T., van der Heijden, G. J., Stijnen, T. & Moons, K. G. Review: A gentle introduction to imputation of missing values. *J. Clin. Epidemiol.* **59**, 1087–1091. <https://doi.org/10.1016/j.jclinepi.2006.01.014> (2006).
22. Furcy, D. & Koenig, S. Limited discrepancy beam search. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence, IJCAI'05*, 125–131 Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, (2005).
23. Zhang, A., Lipton, Z. C., Li, M., & Smola, A. J. 10.8. beam search. In *Dive into Deep Learning, Preview Version* [arXiv:2106.11342](https://arxiv.org/abs/2106.11342) (2023).
24. Kraskov, A., Stögbauer, H. & Grassberger, P. Estimating mutual information. *Phys. Rev. E* **69**, 066138. <https://doi.org/10.1103/PhysRevE.69.066138> (2004).

### Author contributions

R.W., E.S., R.G.M. and A.L. performed research. R.W., E.S., R.G.M. and A.L. wrote the manuscript and prepared the figures. A.L., R.W., E.S., R.G.M., F.D.B, C.T. and F.C. designed the research. E.S., F.D.B., C.T., F.C., F.D. and A.A. collected and curated the data.

### Competing interests

The authors declare no competing interests.

### Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-61334-6>.

**Correspondence** and requests for materials should be addressed to A.L.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024