

SEQUENTIAL DOUBLE CROSS-VALIDATION FOR ASSESSMENT OF ADDED PREDICTIVE ABILITY IN HIGH-DIMENSIONAL OMIC APPLICATIONS

BY MAR RODRÍGUEZ-GIRONDO^{1*}, PERTTU SALO², TOMASZ BURZYKOWSKI³, MARKUS PEROLA², JEANINE HOUWING-DUISTERMAAT^{1,4}, BART MERTENS¹

¹*Department of Medical Statistics and Bioinformatics, Leiden University Medical Center, Leiden, The Netherlands,* ²*National Institute For Health and Welfare, Helsinki, Finland,* ³*Interuniversity Institute for Biostatistics and statistical Bioinformatics (I-BioStat), Hasselt University, Hasselt, Belgium,* ⁴*Department of Statistics, Leeds University, Leeds, United Kingdom.*

Enriching existing predictive models with new biomolecular markers is an important task in the new multi-omic era. Clinical studies increasingly include new sets of omic measurements which may prove their added value in terms of predictive performance. We introduce a two-step approach for the assessment of the added predictive ability of omic predictors, based on sequential double cross-validation and regularized regression models. We propose several performance indices to summarize the two-stage prediction procedure and a permutation test to formally assess the added predictive value of a second omic set of predictors over a primary omic source. The performance of the test is investigated through simulations. We illustrate the new method through the systematic assessment and comparison of the performance of transcriptomics and metabolomics sources in the prediction of body mass index (BMI) using longitudinal data from the Dietary, Lifestyle, and Genetic determinants of Obesity and Metabolic syndrome (DILGOM) study, a population-based cohort from Finland.

1. Introduction. During the past decade, much attention has been devoted to accommodate single high-dimensional sources of molecular data (omics) in the calibration of prediction models for health traits. For example, microarray-based transcriptome profiling and mass spectrometry proteomics have been established as promising omic predictors in oncology [1, 2, 3] and, to lesser extent, in metabolic health [4, 5]. Nowadays, due to technical advances in the field and evolving biological knowledge, novel omic measures, such as NMR proteomics and metabolomics [6, 7] or nano-LCMS and UPLC glycomics [8, 9] are emerging as potentially powerful new biomolecu-

Keywords and phrases: Added predictive ability; double cross-validation; regularized regression; multiple omics sets

lar marker sets. As a result, it is increasingly common for studies to collect a range of omic measurements in the same set of individuals, using different measurement platforms and covering different aspects of human biology. This causes new statistical challenges, including the evaluation of the ability of novel biomolecular markers to improve predictions based on previously established predictive omic sources, often referred as their added or incremental predictive ability [10, 11, 12].

An illustrative example of these new methodological challenges is given by our motivating problem. We have access to data from 248 individuals sampled from the Helsinki area, Finland, within the Dietary, Lifestyle, and Genetic determinants of Obesity and Metabolic syndrome (DILGOM) study [6]. One-hundred-thirty-seven highly correlated NMR serum metabolites and 7380 beads from array-based transcriptional profiling from blood leukocytes were measured at baseline in 2007, together with a large number of clinical and demographic factors which were also measured in 2014, after seven years of follow-up. Our primary goal is the prediction of future body mass index (BMI), using biomolecular markers measured at baseline. More specifically, we would like to compare the predictive ability of the available metabolomics and transcriptomics, and to determine if both data source should be retained in order to improve single-omic source predictions of BMI at seven years of follow-up.

In our setting, it is necessary to both calibrate the predictive model based on each source of omic predictors and assess the incremental predictive ability of a secondary one relative to the first set, using the same set of observations. Hence, in order to avoid overoptimism and provide realistic estimates of performance, it is necessary to control for the re-use of the data, which has already been employed for model fitting within the same observations [13, 14, 15]. This is a very important issue in omic research, where external validation data are hard to obtain. It is well known that biased estimation of model performance due to re-use of the data increases with large number of predictors [16] and omic sets are typically high-dimensional ($n < p$, n sample size and p the number of predictors). Extra difficulties in our setting are the different dimensions (number of features), scales and correlation structure of each omic source, and possible correlation between omic sources induced by shared underlying biological factors.

Evaluating the added predictive ability of new biomarkers regarding classical, low dimensional, settings has been a topic of intense debate in the biostatistical literature in the last years (see, for example, [11, 12, 18, 19] and references therein). Getting meaningful summary measures and valid statistical procedures for testing the added predictive value are difficult tasks,

even when considering the addition of a single additional biomarker in the classical regression context. In particular, widely used testing procedures for improvement in discrimination based on area under the ROC curve (AUC) differences [17] and net reclassification index (NRI) [10] have shown unacceptable false positive rates in recent simulation studies [18, 19]. Overfitting is a big problem when comparing estimated predictions coming from nested regression models fitted to the same dataset. Moreover, the distributional assumptions of the proposed tests seem inappropriate, translating into poor performance of the aforementioned tests even when using independent validation sets [19].

To date, little attention has been paid to the evaluation of the added predictive ability in high-dimensional settings, where the aforementioned problems are larger and new ones appear, such as the simultaneous inclusion in a unique prediction model of predictors sets of very different nature. Tibshirani and Efron [20] have shown that overfitting may dramatically inflate the estimated added predictive ability of omic sources with respect to a low-dimensional set of clinical parameters. To solve this issue, they have proposed to first create a univariate ‘pre-validated’ omic predictor based on cross-validation techniques [13, 14, 21, 22, 23] and incorporate it as a new covariate to the regression with low-dimensional clinical parameters. In a subsequent publication, Hoefling and Tibshirani [24] have shown that standard tests in regression models are biased for pre-validated predictors. As a solution, the authors suggest a permutation test which seems to perform well under independence of clinical and omic sets. Boulesteix and Hothorn [25] have proposed an alternative method for the same setting of enriching clinical models with a high-dimensional set of predictors. In contrast to [20, 24], they first obtain a clinical prediction based on traditional regression techniques. In a second step, the clinical predictor is incorporated as an offset term in a boosting algorithm based on the omic source of predictors. Previous calibration of the clinical prediction is not addressed in the second step and the same permutation strategy as developed by Hoefling and Tibshirani [24] is used to derive p-values.

In this paper, we propose a two-step procedure for the assessment of additive predictive ability regarding two high-dimensional and correlated sources of omic predictors. To the best of our knowledge, no previous work has addressed this problem before. Our approach combines double cross-validation sequential prediction based on regularized regression models and a permutation test to formally assess the added predictive value of a second omic set of predictors over a primary omic source.

2. Methods. Let the observed data be given by $(\mathbf{y}, \mathbf{X}_1, \mathbf{X}_2)$, where $\mathbf{y} = (y_1, \dots, y_n)^\top$ is the continuous outcome measured in n independent individuals and \mathbf{X}_1 and \mathbf{X}_2 are two matrices of dimension $n \times p$ and $n \times q$, respectively, representing two high-dimensional omic predictor sources with p and q features ($p, q > n$). The main goal is to evaluate the incremental or added value of \mathbf{X}_2 beyond \mathbf{X}_1 in order to predict \mathbf{y} in new observations. Our approach is based on comparing the performance of a primary model based only on \mathbf{X}_1 with an extended model based on \mathbf{X}_2 and adjusted by the primary fit based on \mathbf{X}_1 .

2.1. Sequential estimation with two sources of predictors. We propose a two-step procedure based on the replacement of the original (high-dimensional) sources of predictors by their corresponding estimated values of \mathbf{y} based on a single-source-specific prediction model.

In the first step, we build a prediction model for \mathbf{y} based on \mathbf{X}_1 and a given model specification f . Based on the fitted model, the fitted values $\mathbf{p}_1 = \hat{f}(\mathbf{X}_1) = (p_{11}, \dots, p_{1n})^\top$ are estimated. Then, for each individual i , we take the residual $r_i = y_i - p_{1i}$. We consider $\mathbf{r} = (r_1, \dots, r_n)^\top$ as new response and construct a second prediction model based on \mathbf{X}_2 as predictor source:

$$(1) \quad \mathbf{p}_2 = \mathbf{E}(\mathbf{r}|\mathbf{X}_2) = \mathbf{f}(\mathbf{X}_2)$$

This is equivalent to including \mathbf{p}_1 as an offset term (fixed) in the model based on \mathbf{X}_2 for the prediction of the initial outcome \mathbf{y} :

$$(2) \quad \mathbf{p}_2 = \mathbf{E}(\mathbf{y}|\mathbf{X}_2, \mathbf{p}_1) = \mathbf{f}(\mathbf{X}_2) + \mathbf{p}_1$$

Several statistical methods are available to derive prediction models of continuous outcomes in high-dimensional settings. In this work, we focus on regularized linear regression models [16], where $f(\mathbf{X}) = \mathbf{X}\boldsymbol{\beta}$ and the estimation of $\boldsymbol{\beta}$ is conducted by solving $\min_{\boldsymbol{\beta}} (\mathbf{X}\boldsymbol{\beta} - \mathbf{Y})^\top (\mathbf{X}\boldsymbol{\beta} - \mathbf{Y}) + \lambda \text{pen}(\boldsymbol{\beta})$, where $\text{pen}(\boldsymbol{\beta}) = \frac{1-\alpha}{2} \|\boldsymbol{\beta}\|_2^2 + \alpha \|\boldsymbol{\beta}\|_1$. The penalty parameter λ regularizes the $\boldsymbol{\beta}$ coefficients, by shrinking large coefficients in order to control the bias-variance trade-off. The pre-fixed parameter α determines the type of imposed penalization. We consider two widely used penalization types: $\alpha = 0$ (ridge, i.e., ℓ_2 type penalty [26]) and $\alpha = 1$ (lasso, i.e., ℓ_1 penalty [27]). Note that other model building strategies for prediction of continuous outcomes could have been used in this framework, such as the elastic net penalization [28] by setting $\alpha = 0.5$, or boosting methods [40, 29, 41], among others.

2.2. Double cross-validation prediction. The use of a previously estimated quantity (\mathbf{p}_1) in the calibration of a prediction model based on \mathbf{X}_2

(expressions (1) and (2)) requires, in absence of external validation data, the use of resampling techniques to avoid bias in the assessment of the role of \mathbf{p}_1 and \mathbf{p}_2 . We use double cross-validation algorithms [20, 22, 23, 24], consisting of two nested loops. In the inner loop a cross-validated grid-selection is used to determine the optimal prediction rule, i.e., for model selection, while the outer loop is used to estimate the prediction performance. In our setting, the outer loop allows obtaining ‘predictive’-deletion residuals which fully account for the inherent uncertainty of model fitting on the primary source (\mathbf{X}_1), before assessing the added predictive ability of \mathbf{X}_2 , given by \mathbf{p}_2 . In this manner, we avoid bias in estimates of predictive ability which would result from use of a single-cross-validatory approach only. The basic structure of the double cross-validation procedure to estimate unbiased versions of \mathbf{p}_1 and \mathbf{p}_2 is as follows:

Step 1

Input: \mathbf{y}, \mathbf{X}_1

Output: $\mathbf{p}_1 = (p_{11}, \dots, p_{1n})^\top$

1. Randomly split sample S in J mutually exclusive and exhaustive sub-samples of approximately equal size $S_{[1]}, \dots, S_{[J]}$
2. **for** $j \leftarrow 1$ **to** J , **do**
 - a. $S_{[-j]} = S - S_{[j]}$
 - b. Randomly split sample $S_{[-j]}$ in K sub-samples $S_{[-j;1]}, \dots, S_{[-j;K]}$
 - c. **for** $k \leftarrow 1$ **to** K , **do**
 - i. $S_{[-j;-k]} = S_{[-j]} - S_{[-j;k]}$
 - ii. Fit regression model $\mathbf{y} = \hat{f}_{\lambda_l}^{[-j;-k]}(\mathbf{X}_1) + \epsilon$ for a grid of values of shrinkage parameters $\lambda_l, l = 1, \dots, L$ to $S_{[-j;-k]}$
 - iii. Evaluate $\hat{f}_{\lambda_l}^{[-j;-k]}, l = 1, \dots, L$ in the k -th held-out sub-sample $S_{[-j;k]}$ by calculating $\hat{e}_{\lambda_l}^{[k]} = \sum_{(\mathbf{y}, \mathbf{X}_1) \in S_{[-j;k]}} (\mathbf{y} - \hat{f}_{\lambda_l}^{[-j;-k]}(\mathbf{X}_1))^2$
 - iv. **end for**
 - c. Compute overall cross-validation error: $\hat{e}_{\lambda_l}^{[-j]} = \frac{1}{K} \sum_{k=1}^K \hat{e}_{\lambda_l}^{[k]}$, $l = 1, \dots, L$
 - d. Choose $\lambda_{opt} = \min_{l=1, \dots, L} (\hat{e}_{\lambda_l}^{[-j]})$ and calculate predictions of \mathbf{y} in the j -th held-out sub-sample $S_{[j]}$, $\mathbf{p}_1^{[j]} = \hat{f}_{\lambda_{opt}}^{[-j;-k]}(\mathbf{X}_1)$, $(\mathbf{y}, \mathbf{X}_1) \in S_{[j]}$
 - e. **end for**

3. The vector of predictions of \mathbf{y} , $\mathbf{p}_1 = (p_{11}, \dots, p_{1n})^\top$ is obtained by concatenating the J $\mathbf{p}_1^{[j]}$, $j = 1, \dots, J$ vectors, i.e., $\mathbf{p}_1 = (\mathbf{p}_1^{[1]}, \dots, \mathbf{p}_1^{[J]})^\top$

Step 2

Input: \mathbf{y} , \mathbf{p}_1 , \mathbf{X}_2

Output: $\mathbf{p}_2 = (p_{21}, \dots, p_{2n})^\top$

4. Compute double cross-validated residuals $\mathbf{r} = (y_1 - p_{11}, \dots, y_n - p_{1n})$
5. Repeat **Step 1** considering \mathbf{r} as outcome and \mathbf{X}_2 as set of predictors. Note that this is equivalent to obtaining the double cross-validation predictions of \mathbf{y} based on \mathbf{X}_2 considering \mathbf{p}_1 as offset variable in the J fits in **2.c.ii**.

2.3. *Summary measures of predictive accuracy based on double cross-validation.* In order to evaluate the performance of the sequential procedure introduced in Subsection 2.1., we propose three measures of predictive accuracy, denoted by $Q_{\mathbf{X}_1}^2$, $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$, and $Q_{\mathbf{X}_1, \mathbf{X}_2}^2$, based on sum of squares of the double cross-validated predictions \mathbf{p}_1 and \mathbf{p}_2 , obtained following the procedure described in Subsection 2.2. These summary measures can be regarded as high-dimensional equivalents of calibration measurements for continuous outcomes in low-dimensional settings [30, 31], and an extension of previously discussed proposals in the cross-validation literature [14].

Denote by $\text{PRESS}(\mathbf{y}, \mathbf{p}) = \sum_{i=1}^n (y_i - p_i)^2$ the prediction sum of squares based on a vector of predictions \mathbf{p} , obtained according to some arbitrary model f , $\mathbf{p} = (p_1, \dots, p_n)^\top = E(\mathbf{y}|\mathbf{X})$ and by $\text{CVSS}(\mathbf{p}_1, \mathbf{p}_2) = \sum_{i=1}^n (p_{1i} - p_{2i})^2$ the sum of squared differences between two cross-validated vectors of predictions, e.g., $\mathbf{p}_1 = E_{f_1}(\mathbf{y}|\mathbf{X}_1)$, $\mathbf{p}_2 = E_{f_2}(\mathbf{y}|\mathbf{X}_2)$. Let \mathbf{p}_0 be the simplest cross-validated predictor of \mathbf{y} , based on an intercept-only model. The first step of the sequential procedure can be summarized by:

$$(3) \quad Q_{\mathbf{X}_1}^2 = \frac{\text{CVSS}(\mathbf{p}_1, \mathbf{p}_0)}{\text{PRESS}(\mathbf{y}, \mathbf{p}_0)} = \frac{\sum_{j=1}^J \sum_{j \in S_{[j]}} (p_{1j} - \bar{y}^{[-j]})^2}{\sum_{j=1}^J \sum_{j \in S_{[j]}} (y_j - \bar{y}^{[-j]})^2}.$$

Intuitively, $Q_{\mathbf{X}_1}^2$ represents the proportion of the variation of the response \mathbf{y} that is expected to be explained by $f(\mathbf{X}_1)$ in new individuals, re-scaled by the total amount of prediction variation in the response \mathbf{y} . When $\mathbf{p}_1 = \mathbf{p}_0$ (\mathbf{X}_1 as predictive as a null model based on the mean of \mathbf{y}) $Q_{\mathbf{X}_1}^2 = 0$ and $Q_{\mathbf{X}_1}^2 = 1$ if $\mathbf{p}_1 = \mathbf{y}$.

Assuming that $Q_{\mathbf{X}_1}^2 > 0$, the contribution of the second omic source, \mathbf{X}_2 , in the prediction of \mathbf{y} can be summarized by

$$(4) \quad Q_{\mathbf{X}_2|\mathbf{X}_1}^2 = \frac{\text{CVSS}(\mathbf{p}_2, \overline{\mathbf{y} - \mathbf{p}_1})}{\text{PRESS}(\mathbf{y} - \mathbf{p}_1, \overline{\mathbf{y} - \mathbf{p}_1})} = \frac{\sum_{j=1}^J \sum_{j \in S_{[j]}} \left(p_{2j} - \overline{(y_j - p_{1j})}^{[-j]} \right)^2}{\sum_{j=1}^J \sum_{j \in S_{[j]}} \left(y_j - p_{1j} - \overline{(y_j - p_{1j})}^{[-j]} \right)^2}.$$

$Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ accounts for the predictive capacity of $f(\mathbf{X}_2)$, after removing the part of variation in \mathbf{y} that can be attributed to the first source of predictors \mathbf{X}_1 . Its computation relies on the squared difference between \mathbf{p}_2 (the double cross-validated predictions resulting from the second step of the proposed procedure in Subsection 2.2.) and the corresponding residual from the step 1 ($\mathbf{r} = \mathbf{y} - \mathbf{p}_1$) based on \mathbf{X}_1 , re-scaled by the remaining predicted variation on \mathbf{y} after the first step of the procedure. As a result, $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ can be regarded as the expected ability of \mathbf{X}_2 to predict the part of \mathbf{y} , after adjusting for the predictive capacity of $f(\mathbf{X}_1)$ and accounting for all model fitting in the first stage of the assessment. Note that in Step 1 of the sequential procedure J models are fitted, each based on $S_{[-j]}$, providing residuals with expected zero mean (given specification (1)), i.e., $\overline{(y - p_1)}^{[-j]} \approx 0, j = 1, \dots, J$. Hence, $\sum_{j=1}^J \sum_{j \in S_{[j]}} \left(y_j - p_{1j} - \overline{(y - p_1)}^{[-j]} \right)^2 \approx \sum_{i=1}^n (y_i - p_{1i})^2$ and thus

$$(5) \quad Q_{\mathbf{X}_2|\mathbf{X}_1}^2 \approx \frac{\sum_{i=1}^n p_{2i}^2}{\sum_{i=1}^n (y_i - p_{1i})^2}.$$

Finally, we derive a third summarizing measurement of the overall sequential process, $Q_{\mathbf{X}_1, \mathbf{X}_2}^2$, defined as follows:

$$(6) \quad Q_{\mathbf{X}_1, \mathbf{X}_2}^2 = \frac{\text{CVSS}(\mathbf{p}_1 + \mathbf{p}_2, \mathbf{p}_0)}{\text{PRESS}(\mathbf{y}, \mathbf{p}_0)} = \frac{\sum_{j=1}^J \sum_{j \in S_{[j]}} \left(p_{1j} + p_{2j} - \overline{y}^{[-j]} \right)^2}{\sum_{j=1}^J \sum_{j \in S_{[j]}} \left(y_j - \overline{y}^{[-j]} \right)^2}.$$

$Q_{\mathbf{X}_1, \mathbf{X}_2}^2$ represents the total predictive capacity of the overall sequential procedure based on \mathbf{X}_1 and \mathbf{X}_2 , i.e., the combined predictive ability of \mathbf{X}_1 and \mathbf{X}_2 given by $\mathbf{p}_1 + \mathbf{p}_2$.

The three introduced measures jointly summarize the performance of the two omic sources under study and their interplay in order to predict the outcome \mathbf{y} . The three measurements vary between 0 (null predictive ability) and the maximal value of 1. The interpretation of $Q_{\mathbf{X}_1}^2$ is straightforward, as it simply captures the predictive capacity of the firstly evaluated omic

source. Note that the difference between $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ and $Q_{\mathbf{X}_1,\mathbf{X}_2}^2$ relies on the denominator. In general, if \mathbf{X}_1 is informative, the denominator in expression (4) will be smaller than in expression (6). Thus, the residual variation after Step 1 will be smaller than the total initial variation.

The three summary measures are related by the following expression:

$$(7) \quad (1 - Q_{\mathbf{X}_1}^2)(1 - Q_{\mathbf{X}_2|\mathbf{X}_1}^2) \approx (1 - Q_{\mathbf{X}_1,\mathbf{X}_2}^2).$$

Consequently, we can rewrite $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ as follows:

$$(8) \quad Q_{\mathbf{X}_2|\mathbf{X}_1}^2 \approx \frac{Q_{\mathbf{X}_1,\mathbf{X}_2}^2 - Q_{\mathbf{X}_1}^2}{(1 - Q_{\mathbf{X}_1}^2)}.$$

In cases in which $Q_{\mathbf{X}_2|\mathbf{X}_1}^2 = 0$, we get that $Q_{\mathbf{X}_1,\mathbf{X}_2}^2 - Q_{\mathbf{X}_1}^2 = 0$, and vice versa. However, $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ and $Q_{\mathbf{X}_1,\mathbf{X}_2}^2 - Q_{\mathbf{X}_1}^2$ differ when not zero. Specifically, from expression (8), we obtain that $Q_{\mathbf{X}_2|\mathbf{X}_1}^2 \geq Q_{\mathbf{X}_1,\mathbf{X}_2}^2 - Q_{\mathbf{X}_1}^2$. In short, $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ may be regarded as the conditional contribution of \mathbf{X}_2 for the prediction of \mathbf{y} with respect to what may be predicted using \mathbf{X}_1 alone. $Q_{\mathbf{X}_1,\mathbf{X}_2}^2 - Q_{\mathbf{X}_1}^2$ measures the absolute gain in predictive ability from adding \mathbf{X}_2 to \mathbf{X}_1 . Note that a given source \mathbf{X}_2 may present a large conditional $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ but a small absolute $Q_{\mathbf{X}_1,\mathbf{X}_2}^2 - Q_{\mathbf{X}_1}^2$ (if, for example, \mathbf{X}_1 presents high predictive ability itself). Moreover, due to the relation between \mathbf{p}_1 and the resulting vector of predictions after combining \mathbf{X}_1 and \mathbf{X}_2 , $\mathbf{p}_1 + \mathbf{p}_2$, expression (8) implies that $Q_{\mathbf{X}_1,\mathbf{X}_2}^2 \geq Q_{\mathbf{X}_1}^2$. This desirable property may not be fulfilled using alternative combination strategies.

In practice, our sequential procedure relies on the realistic assumption of positive predictive ability of the first source of predictors, \mathbf{X}_1 (one would only be interested in assessing additional or incremental information on top of an informative source itself). Accordingly, we advise to conduct our sequential procedure using \mathbf{X}_1 as primary source only if $Q_{\mathbf{X}_1}^2 > 0$, which is, furthermore, required to derive expression (8).

2.4. Permutation test for assessment of added predictive ability. We introduce a formal test for assessing the added or augmented predictive value of \mathbf{X}_2 over \mathbf{X}_1 to predict \mathbf{y} . We propose a permutation procedure to test the null hypothesis $H_0 : Q_{\mathbf{X}_2|\mathbf{X}_1}^2 = 0$ against the alternative hypothesis $H_1 : Q_{\mathbf{X}_2|\mathbf{X}_1}^2 > 0$. The test is based on permuting the residuals obtained after applying the first step of our two-stage procedure with the data at hand. Our goal is to remove the potential association between \mathbf{X}_2 and \mathbf{y} while preserving the original association between \mathbf{y} and \mathbf{X}_1 . Explicitly, we propose the following algorithm:

Step 1 Calculate the residuals $\mathbf{r} = \mathbf{y} - \mathbf{p}_1$ based on the predictions \mathbf{p}_1 of \mathbf{y} based on \mathbf{X}_1 , obtained in the first step of the procedure presented in Section 2.1.

Step 2 Permute the values of \mathbf{r} , obtaining \mathbf{r}^π and generate values of the response \mathbf{y} under the null hypothesis: $\mathbf{y}^* = \mathbf{p}_1 + \mathbf{r}^\pi$.

Step 3 Repeat the two-stage procedure from Section 2.1. for predicting \mathbf{y}^* and obtain the corresponding $Q_{\mathbf{X}_2|\mathbf{X}_1}^{2*}$.

The procedure is repeated M times and the resulting permutation p-values are obtained as follows:

$$\text{p-value} = \frac{1}{M} \sum_{i=1}^M I(Q_{\mathbf{X}_2|\mathbf{X}_1}^{2*i} > Q_{\mathbf{X}_2|\mathbf{X}_1}^2),$$

where M is the number of permutations, and $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ is the actual observed value with the data at hand. Note that in **Step 2**, we generate a ‘null’ version of the original response \mathbf{y} and then we repeat the overall two-stage procedure, which implies that the ‘null’ residuals used in **Step 3** are not fixed and are, in general, different from \mathbf{r}^π . This is necessary in order to capture all the variability of the two-stage procedure and to correctly generate the null hypothesis of interest. Moreover, the cross-validation nature of the procedure protects against systematic bias of the residuals obtained in Step 3 based on \mathbf{y}^* [see Chapter 7 of 16].

Given the aforementioned relations between $Q_{\mathbf{X}_1}^2$, $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$, and $Q_{\mathbf{X}_1,\mathbf{X}_2}^2$ specified by expression (6), note that $H_0 : Q_{\mathbf{X}_2|\mathbf{X}_1}^2 = 0$ is equivalent to $\tilde{H}_0 : Q_{\mathbf{X}_1,\mathbf{X}_2}^2 - Q_{\mathbf{X}_1}^2 = 0$. This result immediately follows from expression (6), given that $Q_{\mathbf{X}_2|\mathbf{X}_1}^2 = 0$ if and only if $1 - Q_{\mathbf{X}_2|\mathbf{X}_1}^2 = 1$ (assuming $Q_{\mathbf{X}_1}^2 \neq 1$). Hence, both tests are equivalent provided that the distribution under the null hypothesis is generated by the aforementioned permutation procedure, i.e., the p-values, resulting from using $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ as test statistic or $Q_{\mathbf{X}_1,\mathbf{X}_2}^2 - Q_{\mathbf{X}_1}^2$ are approximately the same.

2.5. Software implementation. The proposed method has been implemented in R. Our implementation relies on the R package `glmnet` [32], based on an efficient coordinate descent algorithm which allows to quickly perform **Step 1** and **Step 2** of the double cross-validation prediction introduced in Subsection 2.1. Unfortunately, due to its reliance on resampling, the permutation test proposed in Subsection 2.4 is computational demanding (each permutation iteration takes around 15 seconds for $n = 100$, $p = 1000$, $q = 100$ in a high-end laptop i7 3740 2.7 GHz). Nevertheless, since the permutation procedure is easy to split in M independent realizations of the

same computational procedure, we can use parallel computing to speed up the procedure [33].

3. Simulation study.

3.1. *Simulation setup.* We use matrix singular value decomposition (svd, [34]) of each of the two omic sources to generate common ‘latent’ factors associated with \mathbf{X}_1 , \mathbf{X}_2 , and \mathbf{y} . Common eigenvectors in the svd of \mathbf{X}_1 and \mathbf{X}_2 introduce correlation among the omic sources. We consider different patterns in terms of the conditional association between \mathbf{X}_2 and \mathbf{y} (see Figure 1).

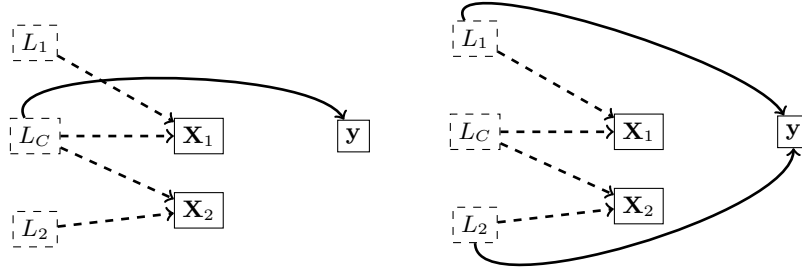


Fig 1: Simulation study. \mathbf{X}_1 and \mathbf{X}_2 are two omic predictors sources and \mathbf{y} is the outcome to be predicted. \mathbf{L}_1 , \mathbf{L}_2 and \mathbf{L}_C are three independent non-observed matrices used to generate \mathbf{X}_1 and \mathbf{X}_2 . Correlation between \mathbf{X}_1 and \mathbf{X}_2 is induced by \mathbf{L}_C . Left: Null case. No independent (of \mathbf{X}_1) association between \mathbf{X}_2 and \mathbf{y} (\mathbf{y} generated as a linear combination of columns of \mathbf{L}_C plus independent noise). Right: Alternative case. Independent (of \mathbf{X}_1) association between \mathbf{X}_2 and \mathbf{y} (\mathbf{y} generated as a linear combination of \mathbf{L}_1 and \mathbf{L}_2 plus independent noise).

The details of the data generation procedure are as follows:

Step 1 Generate $\mathbf{L} \sim N(0, \mathbf{I}_R)$, a matrix of $r = 1, \dots, R$ i.i.d. latent factors of \mathbf{X}_1 and \mathbf{X}_2 .

Step 2 Define Σ_1 ($p \times p$) and Σ_2 ($q \times q$), the correlation matrices of \mathbf{X}_1 and \mathbf{X}_2 , respectively. Following the recent literature on pathway and network analysis of omics data [35], we generated Σ_i , $i = 1, 2$ according to a hub observation model [36, see Figure 2].

Step 3 Draw $\mathbf{X}_i^* \sim N(0, \Sigma_i)$, $i = 1, 2$, and obtain the svd for each of the independent matrices \mathbf{X}_1^* and \mathbf{X}_2^* : $\mathbf{X}_i^* = \mathbf{U}_i^* \mathbf{D}_i \mathbf{V}_i^{*T}$, $i = 1, 2$.

Step 4 Generate the final correlated \mathbf{X}_1 and \mathbf{X}_2 by manipulation of \mathbf{U}^*_{1} and \mathbf{U}^*_{2} , the left eigenvectors matrices from \mathbf{X}^*_{1} and \mathbf{X}^*_{2} , respectively. Specifically, for a certain number (C) of predefined columns C_1 and C_2 , the original submatrices $\mathbf{U}^*_{1C_1}$ and $\mathbf{U}^*_{2C_2}$ (independent) are replaced by C common independent latent factors L_c , $c = 1, \dots, C$ generated in **Step 1**. In this manner, correlation between \mathbf{X}_1 and \mathbf{X}_2 is induced, while the within-omic source correlation structures Σ_1 and Σ_2 are preserved.

Step 5 Simulate the outcome $\mathbf{y} = \mathbf{X}_1\beta_1 + \mathbf{X}_2\beta_2 + \epsilon$, where β_1 and β_2 are vectors of regression coefficients of length p and q , respectively and $\epsilon \sim N(0, 1)$. Since $\mathbf{X}_i = \mathbf{U}_i\mathbf{D}_i\mathbf{V}_i^T$, $i = 1, 2$, we can rewrite $\mathbf{X}_i\beta_i = \mathbf{U}_i\mathbf{D}_i\beta_i^*$ and thus $\beta_i = \mathbf{V}_i\beta_i^*$, where β_i^* represents the association between \mathbf{X}_i and the outcome \mathbf{y} through the orthogonal directions given by \mathbf{U}_i . Consequently, we first generate β_i^* and we then transform it to the predictor space by using $\beta_i = \mathbf{V}_i\beta_i^*$.

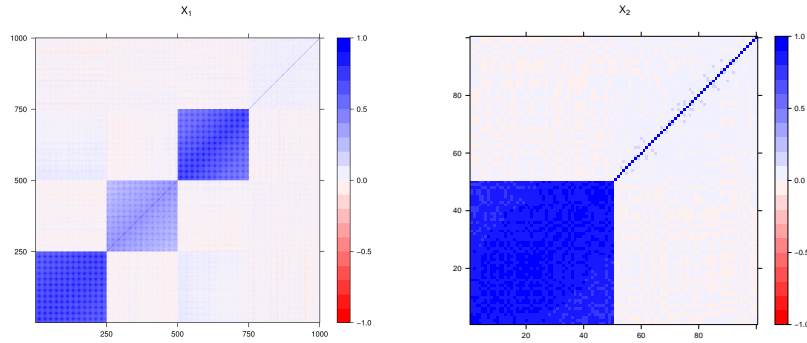


Fig 2: Simulation study. Left: Correlation matrix of \mathbf{X}_1 ($p = 1000$), four groups of 250 features each. Right: Correlation matrix of \mathbf{X}_2 ($q = 100$), two groups of 50 features each.

Simulation 1 ('Null' scenarios): The second omic \mathbf{X}_2 source is non-informative, i.e., $\beta_2^* = \mathbf{0}$, but it is strongly correlated to \mathbf{X}_1 , by imposing common first columns of \mathbf{U}_1 and \mathbf{U}_2 ($\mathbf{U}_{11} = \mathbf{U}_{21}$, the correlation between omic sources is driven through the maximal variance subspace). We considered different assumptions regarding the regression dependence of \mathbf{y} on \mathbf{X}_1 which has an impact on the ability to calibrate prediction rules based on \mathbf{X}_1 for \mathbf{y} . We consider two situations in which the association with \mathbf{y} is unifactorial, in the sense that only one latent factor (one column of \mathbf{U}_1) is associated with \mathbf{y} , and two multi-factorial situations. One of our objectives is to illustrate how changing the complexity of the calibration of a prediction

rule based on \mathbf{X}_1 (by formulating the problem through regression on either larger or smaller variance latent factors) may affect the results. We consider the following ‘null’ scenarios:

(Scenario 1a): $\beta_{1m}^* = 0.01$, $m = 1$; $\beta_{1m}^* = 0$, $m \neq 1$. \mathbf{y} is associated to high-variance subspace of \mathbf{U}_1 , corresponding to the largest eigenvalue of \mathbf{X}_1 .

(Scenario 1b): $\beta_{1m}^* = 0.01$, $m = 6$, $\beta_{1m}^* = 0$, $m \neq 6$. The association with \mathbf{y} relies on a low-variance subspace of \mathbf{U}_1 . Hence, we expect lower values of $Q_{\mathbf{X}_1}^2$, compared to Scenario 1a. **(Scenario 1c):** $\beta_{1m}^* = 0.01$, $m = 1, 2$, $\beta_{1m}^* = 0$ otherwise. In this setting we consider a bifactorial regression, as association with \mathbf{y} is a combination of the effect of the two first eigenvectors of \mathbf{X}_1 .

(Scenario 1d): $\beta_{1m}^* = 0.01$, $m = 1, \dots, 4$, $\beta_{1m}^* = 0$ otherwise. In this setting we consider a multifactorial regression, as association with \mathbf{y} is a combination of the effect of the four first eigenvectors of \mathbf{X}_1 .

Simulation 2 (‘Alternative’ scenarios): \mathbf{X}_2 is associated with \mathbf{y} through latent factors non-shared with \mathbf{X}_1 . The following ‘alternative’ scenarios are investigated:

(Scenario 2a): $\beta_{1m}^* = \beta_{2m}^* = 0.01$, $m = 1$, $\beta_{1m}^* = \beta_{2m}^* = 0$, $m \neq 1$. The eigenvector related to the largest eigenvalue of each source is associated to \mathbf{y} and the association between \mathbf{X}_1 and \mathbf{X}_2 is generated by sharing the second eigenvectors, i.e., by setting $\mathbf{U}_{12} = \mathbf{U}_{22}$. **(Scenario 2b):** $\beta_{1m}^* = 0.01$, $m = 1$, $\beta_{1m}^* = 0$, $m \neq 1$ and $\beta_{2m}^* = 0.01$, $m = 3$, $\beta_{2m}^* = 0$, $m \neq 3$, and the association between \mathbf{X}_1 and \mathbf{X}_2 is generated by setting $\mathbf{U}_{11} = \mathbf{U}_{21}$. **(Scenario 2c):** $\beta_{1m}^* = 0.01$, $m = 6$, $\beta_{1m}^* = 0$, $m \neq 1$ and $\beta_{2m}^* = 0.01$, $m = 3$, $\beta_{2m}^* = 0$, $m \neq 3$, and the association between \mathbf{X}_1 and \mathbf{X}_2 is generated by setting $\mathbf{U}_{11} = \mathbf{U}_{21}$.

Figure 3 shows a Monte Carlo approximation based on a sample of $n = 10,000$ observations of the regression coefficients β_1 and β_2 in the studied simulated scenarios.

In our basic setting, we considered $n = 100$ observations, $p = 1000$ features in \mathbf{X}_1 , and $q = 100$ features in \mathbf{X}_2 . For each scenario, we provide the mean values and standard deviations of $Q_{\mathbf{X}_1}^2$, $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$, and $Q_{\mathbf{X}_1, \mathbf{X}_2}^2$, based on 5-folds double cross-validation, jointly with the rejection proportions for testing $H_0 : Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ along $M = 500$ Monte Carlo trials. We evaluated the permutation test introduced in Subsection 2.3. using $n_{perm} = 200$ permutations. We complemented our empirical evaluations of the proposed sequential double cross-validation procedure by extending our basic simulation setting in two directions. We checked the impact on modifying sample size ($n = 50$) and the complexity of the problem by varying the number of variables considered in the first stage ($p = 4000$).

Additionally, we compared the performance of our procedure based on

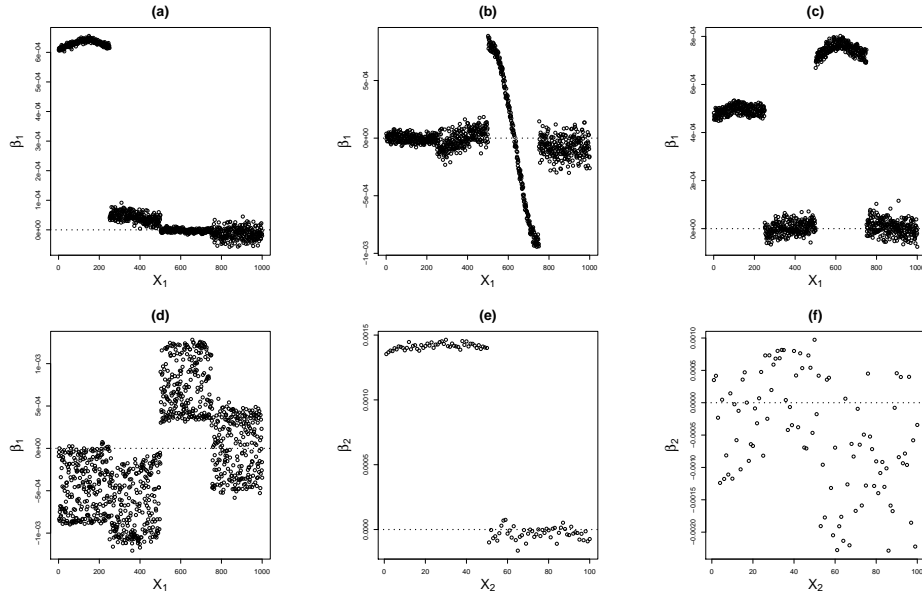


Fig 3: Simulation study. (a)-(d): regression coefficients (elements of β_1 , y-axis) corresponding to each of the p predictors of \mathbf{X}_1 (x-axis). (c)-(d): regression coefficients (elements of β_2 , y-axis) corresponding to each of the q predictors of \mathbf{X}_2 (x-axis). The outcome variable is generated as $\mathbf{y} = \mathbf{X}_1\beta_1 + \mathbf{X}_2\beta_2 + \epsilon$, $\epsilon \sim N(0, 1)$. (e)-(f) provide information about association between \mathbf{y} and \mathbf{X}_1 and (e) and (f) corresponds to the independent association between \mathbf{y} and \mathbf{X}_2 in the alternative scenarios (for the null scenarios 1a-1d the independent association between \mathbf{y} and \mathbf{X}_2 is null). (a) corresponds to scenarios 1a, 2a and 2b, (b) correspond to scenarios 1b and 2c respectively, while (c) and (d) correspond to scenario 1c and 1d. (e) shows the association (β_2) between \mathbf{X}_2 and \mathbf{y} in scenario 2a and (e) shows β_2 for scenarios 2b and 2c.

double-cross validation with two alternative strategies. On the one hand, we provide results based on a two-stage procedure using a single cross-validation loop (cross-validation is used for model choice but predictions and, therefore, the residuals used as outcome in the second stage are directly computed on the complete sample). On the other hand, we check the impact on the results of over-penalization. Specifically, instead of taking λ_{opt} as defined in the inner loop of the double cross-validation procedure presented in Subsection 2.1., we choose a larger value for λ , namely $\lambda_{opt} + 1s.e.(\lambda_{opt})$. The results of these alternative strategies are provided as Supplemental Material but discussed in the main text.

3.2. Simulation results. The results for the sequential double cross-validation procedure (labeled as ‘CV type= CV_D , λ_{opt} ’) are summarized in Tables 1 (ridge regression) and 2 (lasso penalty type).

3.2.1. Ridge regression. For the four ‘null’ scenarios 1a-1d, we expect $Q_{\mathbf{X}_2|\mathbf{X}_1}^2 = 0$ and rejection proportions of H_0 about 0.05. The results of the sequential double cross-validation procedure based on ridge regression are satisfactory in this regard. The top part of Table 1 shows that the estimated $Q_{\mathbf{X}_1}^2$ for scenarios 1a, 1c, and 1d are large and very similar ($Q_{\mathbf{X}_1}^2 \sim 0.90$). As it was expected, the estimated predictive ability of \mathbf{X}_1 is lower in scenario 1b and presents a larger variability, since the association between \mathbf{y} and \mathbf{X}_1 relies on a small variance subspace. In general, for 1a-1d scenarios the estimated $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ is close to zero. However, we observe that the sample size influences the estimated $Q_{\mathbf{X}_1}^2$ and hence, due to the correlation between X_1 and X_2 , also affects the estimation of $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$. We observe systematically lower values of $Q_{\mathbf{X}_1}$ for $n = 50$ than for $n = 100$ in all the studied ‘null’ scenarios. This feature translates to systematically larger values of $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ for $n = 50$ than for $n = 100$. However, the permutation test is able to account for this issue and the level of the test is respected independently of the sample size. Analogously, increasing the number of features of the first source \mathbf{X}_1 (from $p = 1000$ to $p = 4000$) while keeping fixed the number of features of \mathbf{X}_2 ($q = 100$) also affects the estimation of $Q_{\mathbf{X}_1}$ and $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$. In this case, the values of $Q_{\mathbf{X}_1}$ are larger and hence, the values of $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ tend to be closer to zero. Worth noting is that the level of the test is also well respected in this case.

The bottom part of Table 1 shows the results for the alternative scenarios. As desirable, the power increases with sample size for all the three studied alternative scenarios. As it was the case for the ‘null’ scenarios, increasing the sample size tends to lead to better predictive ability of \mathbf{X}_1 . An exception

to this is scenario 2c, where our double cross-validation procedure seems to overfit with $n = 50$. Scenario 2c, unlike scenarios 2a and 2b, is characterized as a ‘difficult’ prediction problem when considering \mathbf{X}_1 (association with \mathbf{y} is driven by a low-variance subspace of \mathbf{X}_1). The greatest power is reached in scenario 2a, in which the independent association between \mathbf{X}_2 and \mathbf{y} is driven through the subspace of maximum variation and the first step of the procedure relies on a relatively ‘easy’ prediction problem. Even if scenarios 2b and 2c are based on the same independent association between \mathbf{X}_2 and \mathbf{y} , the impact of the first source on the power of the test is large. Scenario 2b, in which $Q_{\mathbf{X}_1}^2 = 0.87$ for $n = 100$ reaches a power of 71 %, while the rejection rate reduces to 19% in scenario 2c, corresponding to a more ‘difficult’ prediction problem in the first stage, reflected in a low and unstable $Q_{\mathbf{X}_1}^2$ ($Q_{\mathbf{X}_1}^2 = 0.28$ for $n = 50$ and $Q_{\mathbf{X}_1}^2 = 0.09$ for $n = 100$).

3.2.2. Lasso regression. With regard to the ‘null’ scenarios (top part of Table 2), we observe a good performance for scenarios 1a and 1b, with rejection proportions close to the nominal level. Interestingly, the rejection proportion of the permutation test increases with sample size and the number of features in the first source in scenarios 1c and 1d, which indicates a bad performance of the procedure based on lasso regression in these settings. The reason behind this difference with the ridge-based results is the mis-specification of the lasso with respect to the underlying data-generating mechanism. Lasso regression assumes that the true model is sparse, while scenario 1c and, specially, 1d correspond to non-sparse solutions. These findings illustrate how model mis-specification may result in an improvement of predictions by adding a second source of predictors, not because of independent association to the outcome, but just because of the correlation with the first source of predictors.

The bottom part of Table 2 shows the results for the alternative scenarios. The conclusions are similar to those observed for ridge regression. The power increases with the sample size, and the rejection proportions differ across the three scenarios. However, we observe that ridge outperforms lasso in terms of power, especially for scenarios 2a and 2b.

3.2.3. Alternative procedures. Tables S3 and S4 summarize the results for the two aforementioned alternative strategies in the basic setting ($p = 1000$, $q = 100$) and two sample sizes ($n = 50$, $n = 100$): ‘ CV_D, λ_{1se} ’ corresponds to the strategy in which the sequential double cross-validation is over-shrunk (by taking $\lambda_{opt} + 1s.e.(\lambda_{opt})$ instead of λ_{opt} in the inner cross-validation loop) and ‘ CV_S, λ_{opt} ’ represents the sequential procedure based on one single cross-validation loop (standard residuals as opposed to deletion-based residuals).

TABLE 1
Ridge. Mean estimates (and standard deviation in brackets) of $Q_{\mathbf{x}_1}^2$, $Q_{\mathbf{x}_2|\mathbf{x}_1}^2$, $Q_{\mathbf{x}_1, \mathbf{x}_2}^2$ and rejection proportions of the permutation test based on $Q_{\mathbf{x}_2|\mathbf{x}_1}^2$ along 500 Monte Carlo trials.

| Scenario | n | $Q_{\mathbf{x}_1}^2$ (Step 1) | $Q_{\mathbf{x}_2 \mathbf{x}_1}^2$ (Step 2) | $Q_{\mathbf{x}_1, \mathbf{x}_2}^2$ (Global) | Rej. Prop. |
|----------|-----------------|-------------------------------|--|---|------------|
| 1a | 50 | 0.85 (0.03) | 0.07 (0.09) | 0.89 (0.05) | 0.058 |
| | 100 | 0.88 (0.02) | 0.03 (0.04) | 0.91 (0.03) | 0.068 |
| | 100, $p = 4000$ | 0.94 (0.02) | 0.03 (0.04) | 0.99 (0.01) | 0.056 |
| 1b | 50 | 0.31 (0.07) | 0.04 (0.07) | 0.34 (0.09) | 0.044 |
| | 100 | 0.41 (0.07) | 0.01 (0.02) | 0.42 (0.07) | 0.047 |
| | 100, $p = 4000$ | 0.50 (0.09) | 0.01 (0.03) | 0.72 (0.13) | 0.060 |
| 1c | 50 | 0.86 (0.03) | 0.06 (0.08) | 0.86 (0.04) | 0.060 |
| | 100 | 0.91 (0.02) | 0.02 (0.03) | 0.92 (0.02) | 0.050 |
| | 100, $p = 4000$ | 0.92 (0.05) | 0.02 (0.04) | 0.97 (0.05) | 0.064 |
| 1d | 50 | 0.83 (0.03) | 0.05 (0.08) | 0.84 (0.04) | 0.046 |
| | 100 | 0.86 (0.02) | 0.00 (0.00) | 0.97 (0.01) | 0.062 |
| | 100, $p = 4000$ | 0.88 (0.05) | 0.00 (0.00) | 0.97 (0.05) | 0.044 |
| 2a | 50 | 0.64 (0.13) | 0.50 (0.15) | 0.76 (0.16) | 0.936 |
| | 100 | 0.68 (0.10) | 0.60 (0.12) | 0.91 (0.11) | 0.997 |
| | 100, $p = 4000$ | 0.89 (0.05) | 0.59 (0.08) | 0.93 (0.06) | 0.996 |
| 2b | 50 | 0.84 (0.04) | 0.16 (0.11) | 0.93 (0.04) | 0.236 |
| | 100 | 0.87 (0.03) | 0.11 (0.06) | 0.95 (0.01) | 0.712 |
| | 100, $p = 4000$ | 0.88 (0.02) | 0.10 (0.06) | 0.94 (0.03) | 0.652 |
| 2c | 50 | 0.28 (0.07) | 0.11 (0.11) | 0.36 (0.11) | 0.184 |
| | 100 | 0.09 (0.08) | 0.01 (0.00) | 0.13 (0.16) | 0.186 |
| | 100, $p = 4000$ | 0.16 (0.09) | 0.08 (0.06) | 0.52 (0.09) | 0.526 |

TABLE 2
Lasso. Mean estimates (and standard deviation in brackets) of $Q_{\mathbf{x}_1}^2$, $Q_{\mathbf{x}_2|\mathbf{x}_1}^2$, $Q_{\mathbf{x}_1,\mathbf{x}_2}^2$ and rejection proportions of the permutation test based on $Q_{\mathbf{x}_2|\mathbf{x}_1}^2$ along 500 Monte Carlo trials.

| Scenario | n | $Q_{\mathbf{x}_1}^2$ (Step 1) | $Q_{\mathbf{x}_2 \mathbf{x}_1}^2$ (Step 2) | $Q_{\mathbf{x}_1,\mathbf{x}_2}^2$ (Global) | Rej. Prop. |
|----------|-----------------|-------------------------------|--|--|------------|
| 1a | 50 | 0.79 (0.06) | 0.15 (0.14) | 0.88 (0.05) | 0.058 |
| | 100 | 0.86 (0.03) | 0.05 (0.06) | 0.91 (0.03) | 0.054 |
| | 100, $p = 4000$ | 0.89 (0.02) | 0.10 (0.06) | 0.95 (0.01) | 0.140 |
| 1b | 50 | 0.18 (0.11) | 0.08 (0.12) | 0.27 (0.16) | 0.056 |
| | 100 | 0.33 (0.09) | 0.02 (0.04) | 0.35 (0.09) | 0.056 |
| | 100, $p = 4000$ | 0.44(0.07) | 0.03 (0.05) | 0.45 (0.07) | 0.058 |
| 1c | 50 | 0.80 (0.05) | 0.12 (0.13) | 0.84 (0.05) | 0.054 |
| | 100 | 0.89 (0.02) | 0.04 (0.05) | 0.90 (0.03) | 0.068 |
| | 100, $p = 4000$ | 0.85 (0.28) | 0.13 (0.09) | 0.90 (0.03) | 0.352 |
| 1d | 50 | 0.67 (0.06) | 0.14 (0.12) | 0.72 (0.07) | 0.070 |
| | 100 | 0.83 (0.03) | 0.05 (0.05) | 0.85 (0.03) | 0.094 |
| | 100, $p = 4000$ | 0.73 (0.04) | 0.12 (0.09) | 0.78 (0.04) | 0.360 |
| 2a | 50 | 0.58 (0.09) | 0.44 (0.14) | 0.69 (0.13) | 0.768 |
| | 100 | 0.67 (0.07) | 0.53 (0.09) | 0.90 (0.03) | 1.000 |
| | 100, $p = 4000$ | 0.84 (0.04) | 0.47 (0.09) | 0.84 (0.06) | 1.000 |
| 2b | 50 | 0.77 (0.07) | 0.21 (0.15) | 0.89 (0.06) | 0.106 |
| | 100 | 0.84 (0.04) | 0.12 (0.09) | 0.92 (0.02) | 0.481 |
| | 100, $p = 4000$ | 0.83 (0.02) | 0.18 (0.09) | 0.96 (0.02) | 0.506 |
| 2c | 50 | 0.15 (0.10) | 0.12 (0.13) | 0.28 (0.17) | 0.086 |
| | 100 | 0.27 (0.09) | 0.09 (0.07) | 0.33 (0.11) | 0.351 |
| | 100, $p = 4000$ | 0.07 (0.07) | 0.06 (0.06) | 0.43 (0.07) | 0.227 |

In general, these two alternative strategies provide different estimates for the predictive ability of the two studied sources of predictors. Taking the double-cross validation approach as gold-standard, we observe that the over-shrinkage of the predictions in the first step of the ‘ CV_D, λ_{1se} ’ method provokes an under-estimation of $Q_{\mathbf{X}_1}^2$, while the ‘ CV_S, λ_{opt} ’ provides an over-estimation, specially when the association between outcome and first source of predictors is driven through a low-variance space. Moreover, we observe that the effect of re-using the data is larger for small sample sizes, with systematically larger $Q_{\mathbf{X}_1}^2$ for $n = 50$ than for $n = 100$. However, under the null hypothesis, the introduced bias on the first step for both alternatives does not translate in an inflated type I error. The method ‘ CV_D, λ_{1se} ’ controls the false discovery rate under the null hypothesis in similar fashion than the procedure introduced in Subsection 2.1. With regard to the method based on single cross-validation (‘ CV_S, λ_{opt} ’), its behavior is slightly conservative under the null hypothesis.

For the alternative scenarios, as $Q_{\mathbf{X}_1}^2$, $Q_{\mathbf{X}_2|\mathbf{X}_1}^2$ and $Q_{\mathbf{X}_1, \mathbf{X}_2}^2$, are underestimated by ‘ CV_D, λ_{1se} ’, while ‘ CV_S, λ_{opt} ’ overfits both. Even if power increases with sample size, both methods are systematically less powerful than our proposal, ‘ CV_D, λ_{opt} ’, which makes it the preferable method from both an estimation and testing point of view.

4. Application: DILGOM data. To illustrate the performance of the proposed sequential double cross-validation procedure, and to compare it to the alternative strategies discussed in Section 3, we analyzed data from the DILGOM study, introduced in Section 1. The metabolomic predictor data consists of quantitative information on 137 metabolic measures, mainly composed of measures on different lipid subclasses, but also amino acids, and creatine. The gene expression profiles were derived from Illumina 610-Quad SNParrays (Illumina Inc., San Diego, CA, USA). Initially, 35,419 expression probes were available after quality filtering. In addition to the pre-processing steps described by [6], we conducted a prior filtering approach and removed from our analyses probes with extremely low variation (see [37] for details on the conducted pre-processing). As a result, we retained measures from 7,380 beads for our analyses. The analyzed sample contained $n = 248$ individuals for which both types of omic measurements and the BMI after 7 years of follow-up (mean= 26 kg/m^2 , sd= 5 kg/m^2) were available. We carried out two distinct analyses using the added predictive value assessment approach described in this paper. As a first analysis, we considered the metabolic profile as primary omic source for the prediction of the log-transformed BMI and we evaluated the added predictive value of blood transcriptomics profiles. This

approach is the most relevant in practice, because of both biological and economical reasons. On the one hand, metabolome (which contains, among other, cholesterol measures) is presumably more predictive of BMI than gene expression in blood. On the other hand, NMR technology is typically more affordable [38] than available technologies for transcriptomic profiling, so favoring the NMR source seems a sensible approach in our setting. Nevertheless, to illustrate the properties of our method, we also considered a second analysis in which we reversed the roles of the omic sources, first fitting a model based on gene expression and then evaluating the added predictive value of the metabolome. As in the simulation study, we considered ridge and lasso regression as prediction models, using the same alternative strategies to the sequential double cross-validation procedure presented in section 3 (CV_D, λ_{opt}): CV_D, λ_{1se} , and CV_S, λ_{opt} . The main findings are summarized in Table 3. To check stability of the results, we artificially reduced the sample size of the available DILGOM data and checked the impact on the estimation of the added predictive ability with our sequential double cross-validation approach and its corresponding p-value. We also compared our method with a naive approach, consisting in stacking both metabolites and transcriptomics, and hence ignoring their different origin. The results of these two additional analyses are given as Supplemental Materials.

TABLE 3
Application to DILGOM data. Alternative cross-validation strategies. P-values based on 1000 permutations.

| Method | CV type | Q_{Metab}^2 | $Q_{GE Metab}^2$ | $Q_{Metab,GE}^2$ | p-value |
|--------|-----------------------|---------------|------------------|------------------|---------|
| Ridge | CV_D, λ_{opt} | 0.305 | 0.071 | 0.415 | < 0.001 |
| Ridge | CV_D, λ_{1se} | 0.090 | 0.006 | 0.102 | < 0.001 |
| Ridge | CV_S, λ_{opt} | 0.291 | 0.137 | 0.447 | 0.001 |
| Lasso | CV_D, λ_{opt} | 0.343 | 0.090 | 0.458 | 0.004 |
| Lasso | CV_D, λ_{1se} | 0.073 | 0.004 | 0.083 | 0.018 |
| Lasso | CV_S, λ_{opt} | 0.374 | 0.060 | 0.457 | 0.035 |
| Method | CV type | Q_{GE}^2 | $Q_{Metab GE}^2$ | $Q_{GE,Metab}^2$ | p-value |
| Ridge | CV_D, λ_{opt} | 0.092 | 0.194 | 0.327 | < 0.001 |
| Ridge | CV_D, λ_{1se} | 0.019 | 0.091 | 0.117 | < 0.001 |
| Ridge | CV_S, λ_{opt} | 0.153 | 0.191 | 0.380 | 0.003 |
| Lasso | CV_D, λ_{opt} | 0.277 | 0.102 | 0.453 | < 0.001 |
| Lasso | CV_D, λ_{1se} | 0.043 | 0.122 | 0.204 | < 0.001 |
| Lasso | CV_S, λ_{opt} | 0.372 | 0.012 | 0.409 | 0.490 |

The upper part of Table 3 shows the results from our primary analy-

sis, focused on evaluating the added predictive ability of gene expression with respect to our primary omic set, NMR-metabolomics, in the context of prediction of future BMI. We observe that NMR-metabolomics itself presents a moderately large predictive ability. Both ridge and lasso regression provide very similar Q_{Metab}^2 values which are slightly larger than 30% ($Q_{Metab}^2 = 0.305$ for ridge regression, $Q_{Metab}^2 = 0.343$ for lasso regression). According to the results shown in the third and fourth columns of Table 3, we observe a highly significant positive added predictive ability of gene expression with regard to log-transformed BMI after 7 years of follow-up. The size of such added predictive ability (summarized by $Q_{GE|Metab}^2$) indicates a modest incremental contribution of gene expression, very similar for the two considered regularization methods ($Q_{GE|Metab}^2 = 0.071$ according to ridge regression, $Q_{GE|Metab}^2 = 0.090$ for lasso regression). The final estimated combined predictive ability of metabolites and transcriptomics is slightly larger than 40 % ($Q_{Metab,GE}^2 = 0.415$ according to ridge regression, $Q_{Metab,GE}^2 = 0.458$ for lasso regression). In summary, our main analysis suggests that adding gene expression to an existing model based on NMR-metabolome will potentially lead to a significant improvement in the prediction of BMI at 7 years of follow-up.

The lower part of Table 3 contains the results of our secondary analysis, in which we consider gene expression as primary source and we evaluate the addition of metabolomics. As expected, we observe a better performance of NMR-metabolome than transcriptomics to predict future BMI, for both ridge and lasso regression. In contrast to metabolites, for which the estimated value of Q_{Metab}^2 was only marginally affected by the two considered regularization methods, the type of shrinkage plays a more important role when considering gene expression ($Q_{GE}^2 = 0.092$ with ridge regression, $Q_{GE}^2 = 0.277$ with lasso regression). Accordingly, the estimated added predictive value of metabolites with respect to a primary model based on gene expression is strongly affected by shrinkage type. Namely, metabolome explains around 19% of the remaining variation of the outcome after accounting for the ridge-regression-based prediction calibration using gene expression data, while the equivalent added predictive ability of metabolome drops to 10% if we model gene expression with a lasso penalty in the first step. Interestingly, when considering lasso regression, the impact of the considering transcriptomics or metabolomics as primary source is small, and in fact, the resulting predictive ability after the overall process is similar for both sequential models ($Q_{Metab,GE}^2 = 0.458$, $Q_{GE,Metab}^2 = 0.453$).

The results are supported by our first supplementary analysis consisting of artificially reducing the sample size of the available DILGOM data to

check the stability of the estimated added predictive ability of the considered omic sources, measured through $Q_{Metab|GE}^2$ and $Q_{GE|Metab}^2$. In Table S1 we observe that the results based on a random subsample of the data of $n = 100$ individuals are similar to those obtained with the actual sample, which indicates that the estimated added predictive values shown in Table 3 are reflecting the independent contribution of the secondary source in the prediction of BMI and not the effect of recalibrating the primary source through the second one due to scarcity of data in the first stage. That seems to be the case if we consider a random subsample of the DILGOM data of $n = 50$, when the added predictive ability of the secondary source is presumably overestimated due to the underestimation of the predictive ability of the primary source in the first stage of the procedure.

Our second supplementary analysis also supports our sequential approach. Table S2 shows the estimated values of $Q_{Metab,GE}^2$ for two models in which we simultaneously model both sources of predictors (without and with scaling). These results show that stacking both omic sources is a bad strategy, even with previous scaling of the two sources of predictors to avoid the impact of the different scale of metabolomics and transcriptomics. For both ridge and lasso specifications, we observe that the joint models based on stacking transcriptomics and metabolomics provide lower $Q_{Metab,GE}^2 = Q_{GE,Metab}^2$ values than the initial Q_{Metab}^2 .

With regard to the alternative strategies, Table 3 shows that the use of a single cross-validation approach overestimates the role of gene expression when considered alone ($'CV_S, \lambda_{opt}'$ approach provides a considerably larger Q_{GE}^2 than $'CV_D, \lambda_{opt}'$), while the impact on Q_{Metab}^2 is lower (due to its smaller number of features, 139 metabolites versus more than 7,000 gene expression features). Accordingly, we also observe that the impact of overfitting the residuals is larger when the first source is gene expression. Especially for lasso regression (for which we observe the absolute largest difference between $'CV_S, \lambda_{opt}'$ and $'CV_D, \lambda_{opt}'$ in the first step), we see how the added predictive value of metabolome is obscured by the overfitting of the first source ($Q_{Metab|GE}^2 = 0.012$ clearly not significantly different from 0). With respect to the over-shrinkage of the penalty parameter in the inner loop of the double cross-validation ($'2CV, \lambda_{lse}'$), it leads to a more pessimistic estimation of the summary measures in each of the steps of the procedure, particularly when using ridge regression and the case where the gene expression is the second source to be considered ($Q_{GE|Metab}^2 = 0.014$).

5. Summary and discussion. In this paper, we addressed the problem of evaluating the added predictive ability of a high-dimensional omic

dataset for prediction of continuous outcomes in the context of multiple and possibly correlated omic datasets.

This problem is increasingly relevant in clinical practice. With the emergence of new omic technologies, researchers would like to compare the predictive ability of a primary source of predictors with a secondary source, typically more expensive and noisy, to determine if both sources should be retained in order to improve single-omic predictions, and if that is the case, to provide a predictive rule based on both sources. To achieve this, we proposed a sequential method which consists of considering the vector of residuals based on the primary source of predictors as outcome when fitting a prediction model based on the secondary source of omic predictors. This is equivalent to introducing the vector of individual predictions based on the primary source as an extra covariate with fixed coefficient when fitting a prediction model for the original outcome based on the secondary source of omic predictors. The use of a vector of predictions (which are fitted themselves) in a subsequent prediction model requires cross-validation to account for the uncertainty of calibrating the first source of predictors in the procedure. We have proposed several summary measures, based on double cross-validation predictions. Moreover, we have introduced a permutation test to formally test for added predictive ability of the secondary source.

We have focused on the case in which external validation datasets are not available and, hence, researchers have to fit and compare the predictive performance of different models using the same set of patients (internal validation). This is a common situation in practice. Epidemiological studies, such as the Finrisk (DILGOM), expand their clinical databases by including sequences of novel (omic) biomarker measure sets, which cover different biological processes and which are obtained using different technologies. Access to equivalent data from other studies is typically hard to obtain. Moreover, due to technical and economical reasons, these new sets of markers are typically evaluated on a reduced number of individuals of the study, leading to $n < p$ situations.

Closely connected to the difference between the internal and external validation, is the issue of choice between standard classic (lack-of-fit) residuals and the deletion-prediction residuals employed in this paper. Use of standard residuals may suffice when adding novel predictor data to an established and known (external) risk score, but in general greater caution should be applied. Nevertheless, use of standard residuals from previous model fitting when assessing added-value still seems to be the norm in most analyses. This applies not only to biostatistics [25], but also to related fields such as chemometrics [39]. We have shown through simulations and the analysis of the DILGOM

data that use of deletion-residuals is essential to avoid substantial bias in the assessment of the added value of a secondary predictor set, when added to a primary set using the same patient data.

Another feature of our method is that it does not only account for the predictive capacity of the first predictor set, but also depends on model choice. Hence, our added prediction assessment is not just evaluating the predictive impact of a secondary set of measurements in its own right, but rather the joint impact of choice of predictors, model, and estimation approach.

Our application to real data shows that better predictions can be obtained by adding transcriptomics to a model based on NMR-metabolomics, outperforming single-omic predictions. We have also illustrated the impact of model misspecification in our approach and shown that naive approaches which ignore the different nature, size and scale of the considered source of predictors fail, providing worse results than model based on NMR-metabolomics only.

In our approach, the first source of predictors is prioritized. Several reasons may motivate such an asymmetric approach to combination. On the one hand, available omic sources typically differ in cost and interpretability, and hence researchers may be interested in prioritizing more economic and interpretable sources. This was the case in our real data application based on the DILGOM study. NMR metabolomics measurements are more affordable than transcriptomic profiling, and also more interpretable in the context of BMI prediction. Hence, favoring NMR measurements seems logical in our setting. Similar reasons are typically used to favor classical clinical parameters when evaluating the addition of omic sources to clinical models. One could, for instance, add a novel (biomolecular) marker set, e.g., metabolome, to a set of clinical features such as glucose, blood pressure, and serum cholesterol, which may potentially be correlated to the metabolome markers [20, 24, 25]. On the other hand, our two-stage approach, by fixing the predictive ability of the first source of omic predictors, is robust to differences in scaling and dimensionality between the two sources of predictors.

However, a potential drawback of our sequential approach is its reliance on the order of adding predictors to the model. Alternative non-sequential models (beyond the naive stacking of omic sources) based on simultaneous, but omic-specific, penalization parameters could be used to overcome this limitation. For example, group penalization approaches, based on group lasso [43], sparse group lasso [44], or the recently proposed group ridge [45] could be interesting alternatives to our two-stage approach. Further research is needed to determine how to adapt these methods to our setting in order to correctly deal with differences in scales and dimensionality.

In this work, we have focused on continuous outcome, but our method may

also be immediately extended to other outcome types, particularly to the classification context (binary outcome) and to time-to-event data, for which generalizations of linear regularized regression in high-dimensional settings are available. Summary performance measures as $Q_{\mathbf{x}_1}^2$ and $Q_{\mathbf{x}_1, \mathbf{x}_2}^2$ can be still derived in the binary outcome context. However, given its reliance on the residuals, the extension of the conditional $Q_{\mathbf{x}_2|\mathbf{x}_1}^2$ to the binary and survival contexts is not straightforward. Also, the proposed testing procedure requires major modification when considering more complex responses. Moreover, alternative prediction rules could also be used, beyond linear regularized regression considered here. Boosting algorithms might be an interesting choice worthy of further study in the high-dimensional augmented predictive framework, as an extension of [25], including more complex model specifications, such as non-linear and interaction terms [40, 41, 42]. All these topics require substantial new methodological research and are left as interesting lines of future investigation.

Acknowledgements. Work supported by Grant MIMOmics of the European Union’s Seventh Framework Programme (FP7-Health-F5-2012) number 305280. The DILGOM transcriptomics dataset was funded by Sigrid Jusélius Foundation and Yrjö Jahnsson Foundation. The DILGOM NMR metabolomic dataset was funded by Yrjö Jahnsson Foundation.

References.

- [1] Dudoit, S., Fridlyand, J., and Speed, T. Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American Statistical Association* 2002; **97**: 77-87.
- [2] Rosenwald, A., Wright, G., Chan, W.C., Connors, J.M., Campo, E., et al. The use of molecular profiling to predict survival after chemotherapy for diffuse large-b-cell lymphoma. *New England Journal of Medicine* 2002; **346**: 1937-1947.
- [3] Schwamborn, K. and Caprioli, R.M. Molecular imaging by mass spectroscopy looking beyond classical histology. *Nature Reviews* 2010; **10**: 639-646.
- [4] Apal Sammy, Y.D. and Z.Mohamed. Obesity and genomics: role of technology in unraveling the complex genetic architecture of obesity. *Human Genetics* 2015; **134**: 361-374.
- [5] Jenkinson C.P., Goering, H.H.H., Arya, R., Blangero, J., Duggirala, R., and DeFronzo, R.A.. Transcriptomics in type 2 diabetes: Bridging the gap between genotype and phenotype. *Genomics Data* 2016; **8**: 25-36.
- [6] Inouye, M., Kettunen, J., Soininen, P., Silander, K., Ripatti, S. et al. Metabonomic, transcriptomic, and genomic variation of a population cohort. *Molecular Systems Biology* 2010; **6**: 441.
- [7] Stroeve, J.H., Saccenti, E., Bouwman, J., Dane, A., Strassburg, K., et al. Weight loss predictability by plasma metabolic signatures in adults with obesity and morbid obesity of the DiOGenes study. *Obesity* 2016; **24**: 379-388.
- [8] Zoldos, V., Horvat, T., and Lauc, G. Glycomics meets genomics, epigenomics and

- other high throughput omics for system biology studies. *Current Opinion in Chemical Biology* 2013; **17**: 34–40.
- [9] Theodoratou, E., Thai, K., Agakov, F., Timofeeva, M.N., Stambuk, J., et al. Glycosylation of plasma IgG in colorectal cancer prognosis. *Scientific Reports* 2016; **6**: 28098.
- [10] Pencina M.J., D’Agostino R.B. Sr, D’Agostino R.B. Jr, and Vasan R.S. (2008). Evaluating the added predictive ability of a new marker: from the area under the ROC curve to reclassification and beyond. *Statistics in Medicine* **27**, 207–212.
- [11] Pencina, M.J., D’Agostino, R.B., Pencina, K.M., Janssens, C.J.W., and Greenland, P. (2012). Interpreting incremental value of markers added to risk prediction models *American Journal of Epidemiology* **176**, 473–481.
- [12] Hilden, J., and Gerds, T.A. (2014). A note on the evaluation of novel biomarkers: do not rely on integrated discrimination improvement and net reclassification index. *Statistics in Medicine* **33**, 3405–3414.
- [13] Stone, M. Cross-validatory choice and assessment of statistical predictions (with discussion). *Journal of the Royal Statistical Society. Series B* 1974; **36**: 111–147.
- [14] Jonathan, P., Krzanowski, W.J., and McCarthy, M.V. On the use of cross-validation to assess performance in multivariate prediction. *Statistics and Computing* 2000; **10**: 209–229.
- [15] Varma, S., and Simon, R. Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics* 2006; **7**: 91.
- [16] Hastie, T., Tibshirani, R., and Friedman, J. *Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York: Springer Series in Statistics, 2001.
- [17] DeLong, E.R., DeLong, D.M., and Clarke-Pearson, D.L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics* **44** 837–845.
- [18] Kerr, K.F., Wang, Z., Janes, H., McClelland, R.L., Psaty, B.M., and Pepe, M.S. (2014). Net Reclassification Indices for Evaluating Risk-Prediction Instruments: A Critical Review *Epidemiology* **25**, 114–121.
- [19] Pepe, M.S., Janes, H., and Li, C.I. (2014). Net Risk Reclassification P Values: Valid or Misleading? *Journal of the National Cancer Institute* **106**, dju041 114–121.
- [20] Tibshirani, R.J., and Efron, B. Pre-validation and inference in microarrays. *Statistical Applications in Genetics and Molecular Biology* 2002; **1**: Art. 1.
- [21] Breiman, L. Stacked regressions. *Machine Learning* 1996; **24**: 49–64.
- [22] Mertens, B.J.A., de Noo, M.E., Tollenaar, R.A.E.M. and Deelder, A.M. Mass Spectrometry Proteomic Diagnosis: Enacting the Double CrossValidatory Paradigm. *Journal of Computational Biology* 2006; **13**: 1591–1605.
- [23] Mertens, B.J.A., van der Burgt, Y.E.M., Velstra, B., Mesker, W.E., Tollenaar, R.A.E.M., and Deelder, A.M. On the use of double cross-validation for the combination of proteomic mass spectral data for enhanced diagnosis and prediction. *Statistics and Probability Letters* 2011; **81**: 759–766.
- [24] Hoeffling, H., and Tibshirani, R.J. A study of pre-validation. *The Annals of Applied Statistics* 2008; **2**: 643–664.
- [25] Boulesteix, A.-L., and Hothorn, T. Testing the additional predictive value of high-dimensional molecular data. *BMC Bioinformatics* 2010; **11**: 78.
- [26] Hoerl, A.E., and Kennard, R. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 1970; **12**: 55–67.
- [27] Tibshirani, R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* 1996; **58**: 267–288.
- [28] Zou, H., and Hastie, T. Regularization and Variable Selection via the Elastic Net.

- Journal of the Royal Statistical Society. Series B* 2005; **67**: 301–320.
- [29] Bühlmann, P., and Hothorn, T. Boosting algorithms: regularization, prediction and model fitting. *Statistical Science* 2007; **22**: 477–505.
- [30] Schemper, M. Predictive accuracy and explained variation. *Statistics in Medicine* 2003; **22**: 2299–2308.
- [31] Westerhuis, J.A., Hoefsloot, H.C.J., Smit, S., Vis, D.J., Smilde, A.K., van Velzen, E.J.J., van Duijnhoven, J.P.M., and van Dorsten, F.A. Assessment of PLS-DA cross validation. *Metabolomics* 2008; **4**: 81–89.
- [32] Friedman, J., Hastie, T., and Tibshirani, R. Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software* 2010; **33**:1–22.
- [33] Herlihy, M., and Shavit, N.. *The Art of Multiprocessor Programming (revised edition)*. Elsevier, 2012.
- [34] Jolliffe, I.T. (2008) *Principal component analysis*. Springer-Verlag, New York, second edition.
- [35] Zhang, B. and Horvath, S. A general framework for weighted gene co-expression network analysis. *Statistical Applications in Genetics and Molecular Biology* 2005; **4**: Art. 17, 45pp (electronic).
- [36] Hardin, J., García, S.R. and Golan, D. A method for generating realistic correlation matrices. *The Annals of Applied Statistics* 2013; **7**: 1733–1762.
- [37] Liu, H., D’Andrade, P., Fulmer-Smentek, S., Lorenzi, P., Kohn, K.W., Weinstein, J.N., Pommier, Y., and Reinhold, W.C. mRNA and microRNA expression profiles of the NCI-60 integrated with drug activities *Mol Cancer Ther* 2010; **9**: 1080–1091.
- [38] Soininen, P., Kangas, A.J., Wurtz, P., Tukiainen, T., Tynkkynen, T., Laatikainen, R., Jarvelin, M.R., Kahonen, M., Lehtimäki, T., Viikari, J., Raitakari, O.T., Savolainen, M.J., and Ala-Korpela, M. High-throughput serum NMR metabolomics for cost-effective holistic studies on systemic metabolism. *Analyst* 2009; **134**: 1781–1785.
- [39] Martens, H. and Naes, T. *Multivariate calibration*. Willey & Sons, 1989.
- [40] Tutz, G., and Binder, H. Generalized additive modeling with implicit variable selection by likelihood-based boosting. *Biometrics* 2006; **62**: 961–971.
- [41] Kneib, T., Hothorn, T., and Tutz, G. Variable selection and model choice in geospatial regression models. *Biometrics* 2009; **65**: 626–634.
- [42] Rodríguez-Girondo, M., Kneib, T., Cadarso-Suárez, C., and Abu-Assi, E. Model building in nonproportional hazard regression. *Statistics in Medicine* 2013; **32**: 5301–5314.
- [43] Yuan, M., and Lin, Y. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B* 2006; **68**: 49–67.
- [44] Simon, N., Friedman, J., Hastie, T., and Tibshirani, T. A sparse-group lasso. *Journal of Computational and Graphical Statistics* 2013; **22**: 231–245.
- [45] van de Wiel, M.A., Lien, T.G., Verlaet, W., van Wieringen, W.N., and Wilting, S.M. Better prediction by use of co-data: Adaptive group-regularized ridge regression. *Statistics in Medicine* 2015; **35**: 368–381.