

RESEARCH ARTICLE

Explaining the optimistic performance evaluation of newly proposed methods: A cross-design validation experiment

Christina Nießl^{1,2}  | Sabine Hoffmann^{1,3} | Theresa Ullmann¹ |Anne-Laure Boulesteix¹ 

¹Institute for Medical Information Processing, Biometry and Epidemiology, LMU Munich, Munich, Germany

²Munich Center for Machine Learning (MCML), Munich, Germany

³Department of Statistics, LMU Munich, Munich, Germany

Correspondence

Christina Nießl, Institute for Medical Information Processing, Biometry and Epidemiology, LMU Munich, Marchioninstr. 15, 81377, Munich, Germany.

Email:

cniessl@ibe.med.uni-muenchen.de

Funding information

German Federal Ministry of Education and Research, Grant/Award Number: 01IS18036A; Deutsche Forschungsgemeinschaft, Grant/Award Numbers: BO3139/4-3, BO3139/7-1



This article has earned an open data badge “**Reproducible Research**” for making publicly available the code necessary to reproduce the reported results. The results reported in this article could fully be reproduced.

Abstract

The constant development of new data analysis methods in many fields of research is accompanied by an increasing awareness that these new methods often perform better in their introductory paper than in subsequent comparison studies conducted by other researchers. We attempt to explain this discrepancy by conducting a systematic experiment that we call “cross-design validation of methods”. In the experiment, we select two methods designed for the same data analysis task, reproduce the results shown in each paper, and then reevaluate each method based on the study design (i.e., datasets, competing methods, and evaluation criteria) that was used to show the abilities of the other method. We conduct the experiment for two data analysis tasks, namely cancer subtyping using multiomic data and differential gene expression analysis. Three of the four methods included in the experiment indeed perform worse when they are evaluated on the new study design, which is mainly caused by the different datasets. Apart from illustrating the many degrees of freedom existing in the assessment of a method and their effect on its performance, our experiment suggests that the performance discrepancies between original and subsequent papers may not only be caused by the nonneutrality of the authors proposing the new method but also by differences regarding the level of expertise and field of application. Authors of new methods should thus focus not only on a transparent and extensive evaluation but also on comprehensive method documentation that enables the correct use of their methods in subsequent studies.

KEYWORDS

benchmarking, overoptimism, performance, reproducibility, validation

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Biometrical Journal* published by Wiley-VCH GmbH.

1 | INTRODUCTION

In the literature on data analysis methods, including statistical journals, machine learning journals, and conference proceedings, most articles describe *new* methods, thus contributing to an increasing multitude of potential methods addressing various data analysis problems. It is commonly claimed by the authors proposing these new methods that they perform better than existing ones in some sense. For anecdotal evidence in the context of supervised classification, see Boulesteix et al. (2013). The fact that new methods are typically claimed to be better than existing ones does not necessarily imply that these statements are wrong. In fact, this is what we would expect if we assume continuous scientific progress. However, the recurrent character of these claims, combined with the requirement of journals and reviewers to make these sorts of statements regarding the superiority of the proposed methods, make them somewhat suspicious (Boulesteix et al., 2013; Norel et al., 2011). A recent survey of papers that compare preprocessing methods for a special type of high-throughput molecular data indicates that, at least in this specific context, the paper introducing a method is indeed more optimistic regarding its performance than subsequent papers that are more neutral towards the method (Buchka et al., 2021).

In a different but related approach, several studies demonstrate that it is relatively easy to make a method appear better than it actually is (Jelizarow et al., 2010; Nießl et al., 2022; Pawel et al., 2022; Sonabend et al., 2022; Ullmann et al., 2023). These studies suggest that overoptimistic statements regarding a method's performance may be partly attributed to the nonneutral attitude of the authors, who are naturally interested to present their method in a positive light. More precisely, it is argued that the nonneutrality may translate into a conscious or subconscious optimization of the method and the study design in which it is evaluated (e.g., by selectively reporting the considered datasets or simulation parameters) such that the proposed method shows good performance.

Imagine there are two methods, method 1 and method 2, available to address a specific data analysis task. We set the number of methods to two for the sake of simplicity, but the following arguments can be extended to a setting with more than two methods. Further, imagine the typical situation in which the authors of method 1 and the authors of method 2 both claim that their method performs well. The study designs they use to support their claims are different. We will call them design 1 and design 2, respectively. Following the conjecture discussed in the previous paragraph that study designs used by authors of methods may overfit their methods and vice versa, a natural question is how method 1 would perform when reevaluated using design 2 and how method 2 would perform when reevaluated using design 1. In the present paper, we put this idea into practice by conducting a systematic experiment, which we call “cross-design validation of methods”. More precisely, we consider two exemplary data analysis tasks, namely multiomic data integration for cancer subtyping and differential gene expression analysis, and for each exemplary task we select two papers that propose a new method. For each of these two methods, we reproduce the evaluation shown in the paper that introduced it and then reevaluate it on the design used by the authors of the other paper. In this context of *methodological* research, where data analysis methods are considered as research objects, the study design includes datasets (with a focus on real data for the first task and simulated data for the second task), competing methods, and evaluation criteria.

The goal of this cross-design validation experiment is twofold. First, it allows dissection of the variability of study designs and its impact on the results in the context of methodological research—in a similar way as so-called multianalyst experiments (Silberzahn et al., 2018) do in application fields of statistics. Second, the cross-design validation experiment provides insights into the mechanisms leading to performance discrepancies such as those observed by Buchka et al. (2021) between the original paper (i.e., the paper that introduces the method of interest) and subsequent papers (i.e., papers that propose another method and include the method of interest as competitor or papers that are dedicated to method comparison itself). Importantly, these are *real-world* observations—as opposed to the previous experiments by Jelizarow et al. (2010), Ullmann et al. (2023), and Pawel et al. (2022), mimicking the behavior of fictional researchers who wish to present their method in a favorable light. Finally, our experiment also provides insights regarding the reproducibility of results and the difficulty of performing fair method comparisons as a by-product.

Because the authors of our selected papers made code and data available for the purpose of reproducibility, our experiment can be performed without involving them personally, which considerably simplifies its organization and execution. Moreover, while we could also gain insights from reevaluating the methods of a study design selected by ourselves, the cross-design character of the experiment guarantees a certain degree of neutrality of our comparisons.

The remainder of this paper is structured as follows. The general structure of our cross-design validation experiment is outlined in Section 2. Sections 3 and 4 present the two data analysis tasks, while a discussion of the mechanisms leading to the observed performance differences along with a summary can be found in Section 5. We conclude the paper in Section 6.

2 | PRELIMINARY REMARKS AND DESIGN OF THE EXPERIMENT

2.1 | Terminology

Before describing the experiment in more detail, we briefly clarify the terminology used throughout the paper. Similar to Klau et al. (2020) and Buchka et al. (2021), we define the term *method* not just as the statistical testing or modeling approach, but as the full analysis pipeline potentially including steps such as data normalization. All methods considered in the experiment have several parameters that can be set by the method user (e.g., the maximum number of clusters or the type of multiple testing correction), which we refer to as *method parameters*.

Moreover, we define the *study design* as the combination of all components that contribute to the performance assessment of the method of interest. The study design consists of three main components, namely datasets (real or simulated), competing methods (including their respective method parameters), and evaluation criteria (in our exemplary analysis tasks referring to the evaluation metric and the way the results are aggregated across the real datasets or simulation repetitions). Note that data preprocessing can be seen both as part of the method or part of the data component. In our experiment, we consider preprocessing steps as belonging to the data component if performed for *all* methods and belonging to the method (i.e., the method of interest or the competing methods) otherwise.

2.2 | Selection of the papers

As a preliminary step, we first have to select appropriate papers for both exemplary data analysis tasks that we consider in our experiment, namely cancer subtyping using multiomic data and differential gene expression analysis. Both are applications from the field of biostatistics at the interface with bioinformatics. Apart from the requirement that the paper must introduce a new method, there are two eligibility criteria related to reproducibility: (i) the code to reproduce the results presented in the paper is *publicly* available and can be run without errors, and (ii) the code is written in R, the programming language we are most familiar with. Note that the number of papers to be included in the experiment is not specified in advance and the search for eligible papers is not conducted in a formal or systematic way.

While the restriction to R as a programming language (ii) excludes some papers, the majority of papers fail criterion (i). In many cases, authors only provide the code to use their method (e.g., an R package) but not to reproduce the results shown in the paper. In other cases, the link to the code is broken, the code supposedly included in the supplement cannot be found, or some of the files needed to reproduce the code are missing (e.g., the file containing the empirical data the simulation shown in the paper is based on). Note that we purposely refrain from contacting the authors if the code is not publicly available to make the selection of papers independent of the authors' willingness to respond and provide the code. Although we do not restrict the number of selected papers in advance, the above-mentioned difficulties lead us to stop the search after finding two eligible papers per data analysis task, resulting in $2 \times 2 = 4$ papers included in our experiment.

The conclusion of this search process, although being limited to specific analysis tasks and conducted informally, is that the practice of making code and data openly available is far from being the standard in the methodological literature beyond positive exceptions such as the *Biometrical Journal* (Hofner et al., 2016). The four papers included in our experiment (Nguyen et al., 2019; Osabe et al., 2021; Rappoport & Shamir, 2019; Zhou et al., 2021) should thus be seen as rare positive examples of open research practices in methodological research.

2.3 | Design of the experiment

While all four papers evaluate their respective method extensively in various settings, our experiment includes only the results that (i) are presented as figures or tables and appear in the main paper, that is, excluding the supplement (to keep the experiment feasible) and (ii) compare the method of interest to competing methods (since we can only compare the *relative* performance of a method if the considered papers use different evaluation metrics that do not allow a direct comparison). If the results are based on both real and simulated data, we only consider the results of the data type that is predominantly employed in the paper. In some cases, we exclude more results, which are reported and justified in Sections 3.1 and 4.1.

For each of the four papers, we first compare the results obtained by running the available code to the results presented in the corresponding paper. For this purpose, we use the same R and R package versions that were used by the authors,

TABLE 1 Illustration of the cross-design validation experiment.

	Performance of method 1	Performance of method 2
Study design by authors of method 1	A: Shown in the original paper	B: ?
Study design by authors of method 2	C: ?	D: Shown in the original paper

as far as this information is provided (see Tables S1 and S6 in the Supporting Information). Moreover, we do not modify the code in a way that would change the results, even in cases where we notice discrepancies between the code and the procedure described in the paper (referred to as “design-implementation-gap” by Lohmann et al., 2022, in the context of simulation studies). Exceptions to these rules are explicitly reported in Sections 3.2 and 4.2.

For both data analysis tasks, we then reevaluate each method on the study design used by the authors of the other paper and compare the resulting performances. Our experiment can thus be seen as a “cross-design validation of methods” (see Table 1). As stated above, the study design consists of three main components, namely datasets, competing methods, and evaluation criteria. We also vary these components individually, which allows us to assess their individual impact on the performance of the selected methods. Some challenges arise when reevaluating the methods on the new study design, in particular the choice of method parameters, which we set before viewing the performance results to avoid the risk of favoring one of the methods. Moreover, while we generally adhere to the code used to reproduce the results when “crossing” the designs, some modifications are necessary. Details on how we address these challenges for each data analysis task can be found in Sections 3.2 and 4.2.

The R code and data to reproduce the experiment are openly available at <https://doi.org/10.6084/m9.figshare.20754028>

3 | DATA ANALYSIS TASK I: CANCER SUBTYPING USING MULTIOMIC DATA

The first exemplary data task we consider in our experiment is cancer subtyping through clustering of patients based on multiomic data, an active research field with many newly proposed methods in recent years (see Duan et al., 2021, for an overview). The aim of these methods is to identify clusters (in this context referred to as *subtypes*) with common biological characteristics or clinical phenotypes (e.g., survival time or drug response). This process helps to understand the etiology of the disease and to develop better diagnostic tools and personalized treatment strategies (Duan et al., 2021; Subramanian et al., 2020; Tepeli et al., 2020). Recently developed cancer subtyping methods are usually able to integrate multiple types of high-dimensional molecular data such as genomics, epigenomics, transcriptomics, or proteomics (hence the term *multiomic* data; Subramanian et al., 2020). The two methods selected for our experiment are PINSPlus and NEMO, which were proposed by Nguyen et al. (2019) and Rappoport and Shamir (2019), respectively. Information on where to find the original codes provided by the authors is listed in our code documentation. We will abbreviate Nguyen et al. (2019) and Rappoport and Shamir (2019) by N19 and R19.

3.1 | Study design in the original papers

In the following, we outline and compare the study designs that are used to assess the performance of PINSPlus and NEMO in their respective original papers and that meet the inclusion criteria of our experiment (see Table 2 for an overview). We also report the authors’ justifications for the design choices. For this purpose, we will also refer to T. Nguyen et al. (2017), which propose PINS, the predecessor method of PINSPlus, and to Rappoport and Shamir (2018), a benchmark study intended as neutral that has been previously conducted by the authors of NEMO. All results of NEMO’s competing methods originate from this benchmark study, that is, the results of NEMO were simply added to the results of the previously published benchmark study. Since both R19 and N19 mainly use real datasets to evaluate their methods, we do not further consider the simulation results presented by R19.

Data Both R19 and N19 use datasets from The Cancer Genome Atlas Research Network (TCGA; <https://www.cancer.gov/tcga>), where each dataset corresponds to a different cancer type (e.g., kidney renal clear cell carcinoma or acute myeloid leukemia). The two author teams also consider the same three types of omic data (gene expression, methylation, miRNA expression) but use different numbers of datasets (34 in N19 vs. 10 in R19). Neither N19 nor R19 explicitly comments on the number of datasets and the selected cancer types, although 34 seems to be close to the maximum number of available datasets for the three considered types of omic data at the time of publication. Moreover, neither N19 nor R19 discusses

TABLE 2 Overview of the study design components used for performance assessment of PINSPlus and NEMO.

Study design component		PINSPlus (Nguyen et al., 2019)	NEMO (Rappoport & Shamir, 2019)
Datasets	Number and type	<ul style="list-style-type: none"> 34 TCGA datasets (gene expression, methylation, miRNA expression) *2 METABRIC datasets (gene expression, CNV) 	<ul style="list-style-type: none"> 10 TCGA datasets (gene expression, methylation and miRNA expression) *Partial TCGA datasets
	Preprocessing (all methods)	See Table S2	See Table S2
Competing methods	Number and type	3: SNF, iCluster+, Consensus Clustering	9: PINS, SNF, iClusterBayes, <i>k</i> -means, spectral clustering, MCCA, LRAcluster, *rMKL-LPP, *MultiNMF
	Preprocessing (method-specific)	See Table S2	See Table S2
	Other method parameters	<ul style="list-style-type: none"> SNF: <i>alpha</i> = 0.5, <i>no. iterations</i> = 10, <i>number of clusters</i> = estimated according to eigen-gaps, <i>maximum number of clusters</i> = 5, <i>number of neighbors</i> = 20 Other methods: see the original paper 	<ul style="list-style-type: none"> SNF: <i>alpha</i> = 0.5, <i>number of iterations</i> = 30, <i>number of clusters</i> = estimated according to rotation cost, <i>maximum number of clusters</i> = 15, <i>number of neighbors</i> = number of samples/10 Other methods: see the original paper
Evaluation criteria	Metric	<ul style="list-style-type: none"> Survival: logrank test 	<ul style="list-style-type: none"> Survival: permutation-based logrank test Clinical: permutation-based /Kruskal–Wallis test (discrete/continuous) for up to six clinical variables *Runtime *Number of clusters
	Aggregation	<ul style="list-style-type: none"> Number of datasets with significant and most significant logrank <i>p</i>-value 	<ul style="list-style-type: none"> Number of datasets with significant logrank <i>p</i>-value Number of datasets with at least one enriched clinical variable Mean $-\log_{10}$ logrank <i>p</i>-value Mean number of enriched clinical variables

Note: Included are only components (i) for which the corresponding results are presented as figures or tables in the main paper (i.e., not in the supplement), (ii) that compare the method of interest to other competing methods, and (iii) that correspond to the performance assessment based on real data. In addition, some components are not included in the experiment, which are indicated by asterisks (*). Competing methods and evaluation criteria for datasets not included in the experiment are not shown.

their choice of omic data types, which seems to be a general issue in papers proposing new cancer subtyping methods, as criticized by Duan et al. (2021).

Although the 10 cancer types included by R19 are also considered in N19, the corresponding datasets have different numbers of patients and omic variables. This is mainly caused by the different preprocessing steps applied by N19 and R19 (see Supporting Information Section A.2 for details). In addition, the two papers probably also use different dataset versions (it is not possible to identify the data version used by N19).

Note that N19 also considers two breast cancer datasets that do not originate from TCGA and exhibit different omic types. However, we exclude them from our experiment since some evaluation criteria of R19 require six clinical variables (see below), which are either not available or cannot be clearly identified for these two datasets. Moreover, we do not include the partial datasets (i.e., datasets where some patients do not have any measurements for one or more omic data types) used in R19 to demonstrate NEMO's ability to analyze this type of data. This is because PINSPlus assumes complete data and would require potentially suboptimal solutions such as imputation.

Competing methods R19 and N19 use different numbers and types of competing methods to assess the relative performance of their proposed methods. While R19 uses nine competing methods, N19 only considers three methods. The only method that is included in both papers is similarity network fusion (SNF; Wang et al., 2014). The difference in the number of competing methods is not surprising given that the performance evaluation of NEMO is, in contrast to PINSPlus, based on a benchmark study with a focus on method comparison itself (Rappoport & Shamir, 2018). Such studies typically aim to compare as many methods as possible to generate comprehensive guidelines for method users. Interestingly, R19 includes PINS, the predecessor method of PINSPlus, as a competing method. PINSPlus itself is not included since it did not exist yet when Rappoport and Shamir (2018) conducted their benchmark study. Concerning the choice of competing methods, Rappoport and Shamir (2018) report that they aim to represent diverse multiomic clustering approaches, and that within each approach they choose widely used methods with available software and clear usage guidelines. N19 refer to their selected competing methods as established subtyping methods.

Regarding the parameter selection of the competing methods, NEMO's authors state in Rappoport and Shamir (2018) that they choose the method parameters following the guidelines given by the authors of the respective method (which involves performing a parameter search if suggested) and construct parameter selection methods by themselves if there are no available guidelines. N19 does not have a comparable statement except for the number of clusters for the method consensus clustering (Monti et al., 2003), which, as stated in T. Nguyen et al. (2017), is determined as suggested by Monti et al. (2003). For SNF, the only method that is considered as a competing method for both PINSPlus and NEMO, N19 and R19 both normalize the omic variables to have a mean of 0 and a standard deviation of 1 (which, as stated in Section 2.1, we consider as a method parameter since it is not applied for all methods in both papers). However, they choose different values for the number of neighbors (20 vs. number of samples/10), the number of iterations (10 vs. 30), the number of clusters (estimate according to eigen-gaps vs. rotation cost), and the maximum number of considered clusters (5 vs. 15). See Table S2 for the method-specific preprocessing steps as well as N19 and R19 for all other parameters of the remaining methods.

Note that we have to exclude two competing methods (rMKL-LPP and MultiNMF) considered by R19 from the experiment since we are not able to run them (see Supporting Information Section A.3 for details).

Evaluation criteria With regard to the evaluation criteria, N19 focuses on the methods' ability to identify clusters with significant survival differences using the logrank test. Note that in this context, the logrank test is equivalent to performing a Cox regression (which is the term used by N19), but we will refer to it as the logrank test since this seems to be the more commonly used term in cancer subtyping methodology. T. Nguyen et al. (2017) note that the same logrank test was also used by the authors proposing SNF (Wang et al., 2014), which can be seen as a justification for their choice. For each method, N19 highlights the datasets with significant (i.e., $p < 0.05$), and most significant (i.e., the smallest significant p -value across all methods) p -values by color.

In R19, the assessment of significant survival differences is also based on the logrank test. In addition, the authors assess "clinical enrichment" by testing the association between the identified clusters and six clinical variables (gender, progression of the tumor, cancer in lymph nodes, metastases, total progression, and age at initial diagnosis), although not all variables are available in each clinical dataset. R19 employs the χ^2 -values using a permutation procedure, arguing that in the cancer subtyping context, the χ^2 -values, the number of datasets with at least one enriched clinical variable, the mean $-\log_{10}$ logrank p -value, and the mean number of enriched clinical variables per dataset. R19 thus considers four evaluation criteria regarding survival and clinical enrichment. Note that one of these criteria (the number of datasets with significant logrank p -values) is very similar to the criterion used by N19 (the number of datasets with [most] significant logrank p -values), the only difference being the estimation of the p -value (approximation-based vs. permutation-based) and the inclusion of the number of datasets with the most significant p -values as a second-order ranking criterion in N19.

In addition to analyzing survival differences and clinical enrichment, R19 also reports the number of clusters and the runtime of each method. However, we do not consider these criteria in our experiment since the number of clusters has no clear optimal value and runtime is not comparable due to different computational resources.

3.2 | Challenges when conducting the experiment

Reproducibility The results presented in N19 are fully reproducible, except for one p -value of iCluster+. In contrast, the results presented in R19 cannot be exactly reproduced. Besides the two methods that cannot be run at all, the performance results of the remaining methods are slightly different compared to the original paper, especially for the clinical enrichment criteria (the difference between original and reproduced results with regard to NEMO's performance is reported

in Section 3.3). Interestingly, 76 of the 80 clustering solutions (8 methods \times 10 datasets) are equal to the clustering solutions provided by Rappoport and Shamir (2018), with two of the remaining four solutions only differing in one and three individuals, respectively. This means that the reproducibility problems (also observed for some of the 76 settings yielding identical clustering solutions) might be caused by the permutation tests. Moreover, the provided code is probably not the exact code used by R19, as indicated by the fact that R19 refers to Rappoport and Shamir (2018) for the code to reproduce the results, but also mention that the implementations for MCCA (sparse multiple canonical correlation analysis), LRAcluster, and k -means were slightly changed compared to Rappoport and Shamir (2018).

When reproducing the results, we do not modify the code provided by the authors in a way that would change the results and attempt to use the same R and R package versions as in the original papers (see Table S1 in the Supporting Information). However, we have to set a different number of cores in some settings and use a different R version for running the permutation tests by R19 due to different computational resources (see our code documentation for details), which may contribute to the reproducibility issues.

Crossing the designs Evaluating the performance of PINSPlus and NEMO using each other's datasets, competing methods, and evaluation criteria poses a number of challenges, the most important one being the choice of parameters both for the two methods of interest, PINSPlus and NEMO, and the competing methods. Whenever a method is applied to a new (set of) dataset(s), the method user needs to carefully select its parameters or a corresponding parameter selection method, which of course also applies to our experiment. Since both N19 and R19 use the same three types of omic data from the same source (TCGA), we set the parameters of PINSPlus and NEMO as in their respective original paper, which corresponds to their default parameter setting. Note that we also do not change the range of possible values for the number of clusters, a parameter that can be specified for both methods and is set to $\{2, 3, 4, 5\}$ for PINSPlus and to $\{2, 3, \dots, 14, 15\}$ for NEMO. We also attempt to use the same parameters for the competing methods when applying them to the new datasets. However, for two competing methods of N19 (iCluster+ and Consensus Clustering), the optimal number of clusters has to be selected by the user according to plots generated by the method when run on a specific dataset. When applying these two methods on the datasets by R19, we thus have to manually choose the optimal number of clusters for every dataset, and although we try to imitate the decisions of N19 on their datasets, a clear determination is not always possible (an issue that is also noted by Duan et al., 2021). Moreover, some refinements regarding the method-specific preprocessing steps are necessary for two competing methods of R19 (see Section A.2 in the Supporting Information).

In addition to the choice of method parameters, some challenges arise when applying the evaluation criteria by R19 to the datasets by N19. Specifically, the logrank permutation test by R19 does not converge for some methods on two datasets by N19, resulting in a p -value of 0 in 15 method-data combinations. In these cases, we use the approximation-based logrank p -values. Moreover, clustering solutions resulting from the dataset UCS (N19) are not tested for clinical enrichment (R19) since it only includes one of the six clinical variables ("gender") with only one value ("female").

3.3 | Results

Performance based on the original study design Figure 1A and 1D shows the reproduced performance results of PINSPlus and NEMO based on their original study design. Note that the representation in these figures slightly differs from the original papers to achieve a comprehensive and yet clear summary of the results. The most important difference is that the papers also report the individual performance results for each dataset (we provide the individual performance results in Tables S4 and S5 in the Supporting Information).

When evaluated based on its original design, PINSPlus seems to be clearly superior to the three competing methods. It has the most significant p -values ($p < 0.05$) regarding survival, with 21 of the 25 significant p -values being the smallest across all methods. NEMO also shows good performance in its original study design, although its performance is not as clearly superior to the competing methods as the performance of PINSPlus. It achieves the highest numbers of datasets with significantly different survival and at least one enriched clinical variable (although there are two competing methods that achieve the same number of datasets with clinical enrichment). Moreover, none of the competing methods achieves both a higher mean $-\log_{10}$ logrank p -value and a higher mean number of enriched clinical variables. Only MCCA obtains a higher mean $-\log_{10}$ logrank p -value than NEMO but has a lower mean number of enriched clinical variables. Note that despite the reproducibility issues, both the absolute (i.e., the values of the four evaluation criteria considered by R19) and the relative performance of NEMO (i.e., when comparing these values to the competing methods) correspond to the results shown in the original paper. The only difference affecting the relative performance of NEMO is that in the original paper,

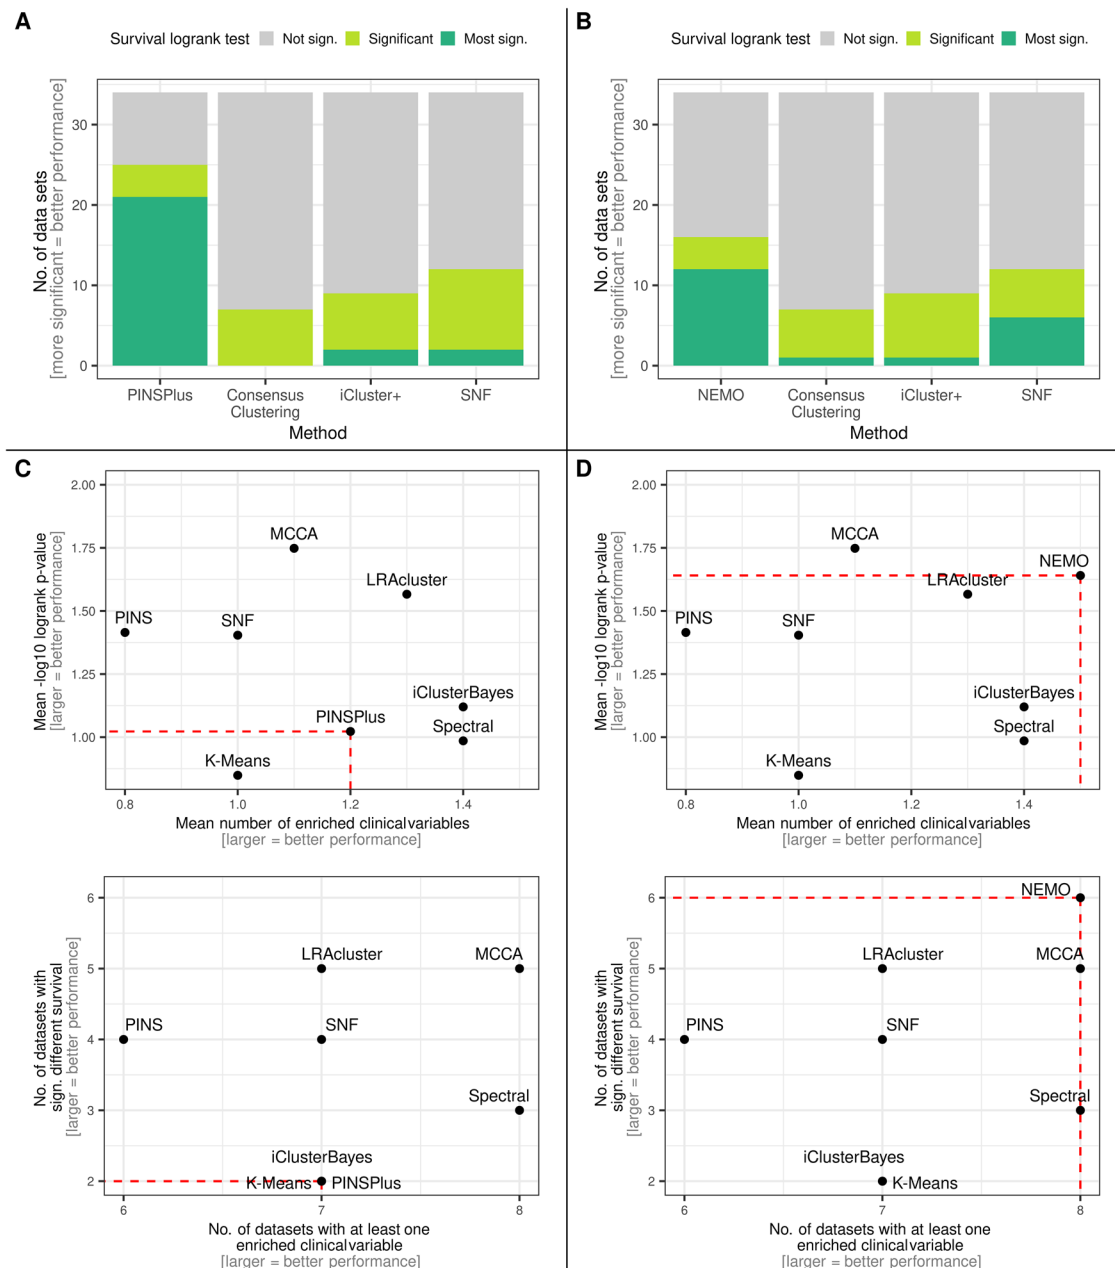


FIGURE 1 Results of the cross-design validation experiment for the cancer subtyping example, where the subfigures A–D correspond to the cells of Table 1. (A) Performance of PINSPPlus based on the design by Nguyen et al. (2019) (= original design). (B) Performance of NEMO based on the design by Nguyen et al. (2019) (= crossed design). (C) Performance of PINSPPlus based on the design by Rappoport and Shamir (2019) (= crossed design). (D) Performance of NEMO based on the design by Rappoport and Shamir (2019) (= original design).

one of the two methods that could not be reproduced (rMKL-LPP), achieves a higher mean number of enriched clinical variables than NEMO but a lower mean $-\log_{10}$ logrank p -value.

Performance based on the crossed design The performance results of NEMO and PINSPPlus based on each others' study design (i.e., datasets, competing methods, and evaluation criteria) are presented in Figure 1B and 1C, respectively. In the study design of R19, PINSPPlus does not outperform the competing methods. It is only the fourth and sixth best method with regard to the mean number of enriched clinical variables and mean $-\log_{10}$ logrank p -value, respectively. It belongs to the three worst methods with regard to the number of datasets with significantly different survival and only outperforms PINS, its predecessor method, with regard to the number of datasets with at least one enriched clinical variable. In contrast, NEMO still outperforms the competing methods in the design by N19, although its superiority is not as pronounced as

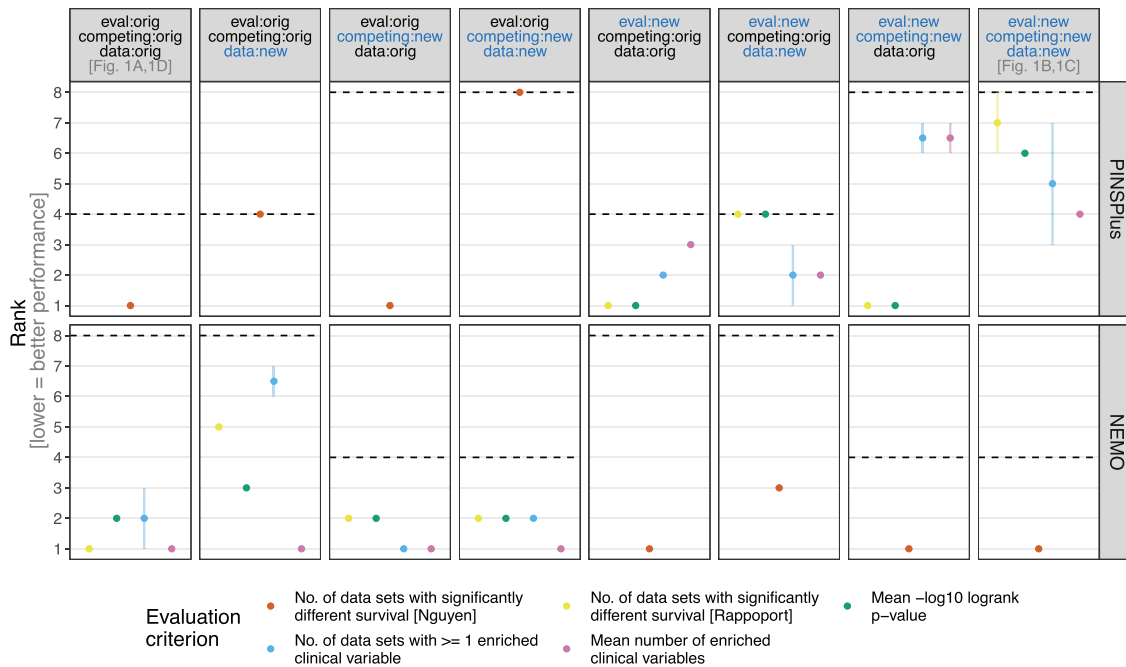


FIGURE 2 Performance ranks of the two cancer subtyping methods PINSPPlus and NEMO based on datasets, competing methods, and evaluation criteria that either correspond to the original (PINSPPlus: Nguyen et al., 2019; NEMO: Rappoport & Shamir, 2019) or crossed design (PINSPPlus: Rappoport & Shamir, 2019; NEMO: Nguyen et al., 2019). Each panel represents the performance rank(s) of PINSPPlus or NEMO for one combination of datasets, competing methods and evaluation criteria. If more than one method achieves the same value for a certain criterion, the point represents the average rank and the line indicates the rank range based on minimum and maximum rank. The dashed lines correspond to the number of compared methods, that is, the highest possible rank.

for PINSPPlus in the same design (PINSPPlus achieves 25 significant p -values while NEMO only achieves 16 for the same 34 datasets).

We also analyze the performance of PINSPPlus and NEMO when datasets, competing methods, and evaluation criteria are varied individually. Figure 2 shows the resulting ranks of PINSPPlus and NEMO for all eight combinations of the three components, where each component can either be set to the original or the crossed version ($2^3 = 8$). For each criterion (one by N19 and four by R19), a rank of 1 corresponds to the best method. If more than one method achieves the same value for a certain criterion, the minimum, maximum, and average ranks are reported. As can be seen from Figure 2, the ranks of PINSPPlus and NEMO generally vary for each combination of datasets, competing methods, and evaluation criterion. Apart from its original design, PINSPPlus achieves rank 1 for the evaluation criteria related to survival (i.e., number of [most] significant p -values and mean $-\log_{10}$ logrank p -value) in all combinations where the datasets by N19 are used. However, PINSPPlus belongs to the worst performing methods according to survival when applied to the datasets by R19. As mentioned in Section 3.1, the 10 datasets corresponding to different cancer types that are used by R19 are also included in N19. Interestingly, PINSPPlus achieves a significant p -value for nine of these 10 datasets in N19, indicating that the difference in performance for these datasets is mainly due to the different preprocessing steps. With regard to the clinical evaluation criteria, PINSPPlus seems to have average performance, neither clearly performing better nor worse than the other methods.

In comparison to PINSPPlus, the ranks of NEMO are more robust across the different study designs. For six of eight study designs, it achieves rank 1 or 2 for all evaluation criteria (if the minimum or average rank is considered). The only study design where NEMO's performance is considerably worse for two evaluation criteria is the design where only the datasets are taken from N19 while evaluation criteria and competing methods correspond to the original paper. Moreover, it can be noted that while the slightly different calculation of the number of datasets with significant logrank p -values in N19 and R19 does not have an impact on the ranks of PINSPPlus, NEMO tends to achieve better ranks for the version of N19. For example, it achieves rank 1 instead of 2, for settings where data and competing methods are by N19. A comparison of approximation-based and permutation-based p -values for all methods and datasets can be found in the Supporting Information file (Figure S1), showing that the approximation-based p -values are indeed generally smaller. The Supporting

Information file also provides a comparison of the two different parameter settings of SNF that are specified by N19 and R19 (Figure S2), which reveals a considerable but nonsystematic performance difference between the two implementations.

4 | DATA ANALYSIS TASK II: DIFFERENTIAL GENE EXPRESSION ANALYSIS

The second data task we consider in our experiment is differential gene expression analysis, which aims at identifying genes that show differences in their expression levels between two or more conditions (Soneson & Delorenzi, 2013). Of the many methods that have been proposed for this task (Seyednasrollah et al., 2013), the more recent ones usually expect RNA-Seq data as input, which means that gene expression is measured as nonnegative counts (Rigaiil et al., 2018). The two methods for differential expression analysis included in the experiment are SFMEB (scaling-free minimum enclosing ball) and MBCdeg (derived from MBCcluster.Seq, a model-based clustering algorithm for RNA-Seq data), which have recently been proposed by Zhou et al. (2021) and Osabe et al. (2021) and require RNA-Seq data as input. As stated in Section 2, these papers are selected because they make the code to reproduce the results openly available (information on where the code can be found is reported in our code documentation). We will abbreviate them by Z21 and O21 in the following.

4.1 | Study design in the original papers

In this section, we review the datasets, competing methods, and evaluation criteria that are used to assess the performance of SFMEB and MBCdeg in their respective original paper and that meet the inclusion criteria of our experiment (see Table 3 for an overview). We also report the justifications for the design choices provided by the authors. Since Z21 and O21 primarily use simulated data to evaluate their methods, we do not further consider their real data analyses.

Data Both Z21 and O21 generate simulated count data representing RNA-Seq read counts of p genes in $2 \times n_{obs}$ samples from two groups. O21 also simulates count data from three groups, but we exclude these settings from the experiment because SFMEB does not seem to be intended for this type of data (all evaluations in the original paper by Z21 are based on two-group data). The simulation framework of Z21 and O21 is based on different code implementations (code by Robinson & Oshlack, 2010, and `compcodeR` R package, Soneson, 2014 vs. TCC R package, Sun et al., 2013) as well as different distributions to generate the count data (Poisson and negative binomial distribution vs. only negative binomial distribution). Moreover, the two papers choose different numbers of simulation repetitions (20 vs. {50,100}), different sample sizes per group ({1,2,5,8} vs. 3), and different numbers of genes ({15,000,..., 29,800} vs. 10,000).

The simulations also differ with respect to the characteristics of the differentially expressed (DE) genes. In contrast to O21, the DE genes in Z21 include uniquely expressed (UE) genes (u_1, u_2) that have zero counts in groups 1 or 2, respectively. Moreover, Z21 and O21 consider different proportions of DE genes ({0.3,...,0.7} excluding UE genes vs. {0.05,...,0.75}), different \log_2 fold-changes between the groups (i.e., the true \log_2 ratio of expression change; ≥ 2 vs. 2), and different proportions of upregulated genes (i.e., genes having higher expression) in group 1 ({0.6,...,1} vs. {0.5,...,1}).

In contrast to O21, Z21 applies prefiltering of the genes (e.g., filtering of genes with mean count ≤ 2) for all methods, although some of their considered methods additionally filter genes internally. Moreover, in some settings, Z21 considers heterogeneous data composed of two datasets with different simulation parameters (\log_2 fold-change, number of genes, etc.). In the results included in the experiment, O21 only varies the proportion of DE genes and the proportion of upregulated genes, but in a fully factorial manner which results in $6 \times 4 = 24$ simulation settings. However, it should be noted that O21 also varies other parameters (e.g., the \log_2 fold-change) in settings not considered in our experiment since they did not meet the inclusion criteria (e.g., because the corresponding figures are shown in the supplement). In the simulation settings by Z21 that are included in our experiment (15 settings in total), more parameters are varied, but not in a fully factorial manner. More specifically, the 15 included settings originate from five “studies” (each consisting of three settings) with different simulation parameters. Within each study, one simulation parameter is varied (see Table 3).

Understandably, neither Z21 nor O21 provides a justification for every single simulation parameter but often refers to similar parameter values observed in real data. Regarding the choice of the number of simulation repetitions, however, neither of the two papers provides a justification. As criticized by Morris et al. (2019), this seems to be a general issue in papers presenting simulation studies.

Competing methods Z21 compares SFMEB with five competing methods they consider as widely used. Two of these methods are referred to as edgeR and DESeq (Anders & Huber, 2010; Robinson et al., 2009; see below for more details), which closely corresponds to the methods selected by O21 (edgeR and DESeq2; Love et al., 2014). In addition to these

TABLE 3 Overview of the study design components used for performance assessment of SFMEB and MBCdeg.

Study design component	SFMEB (Zhou et al., 2021)	MBCdeg (Osabe et al., 2021)
Datasets	Two conditions	Two conditions, *Three conditions
Number of conditions	15	24
Number of settings	Study 1 Study 2 Study 3 Study 4 Study 5	Fig. 1 in O21 Fig. 3 in O21
Setting names	Poisson	Negative binomial
Distribution	Robinson and Oshlack (2010)	Negative binomial
Code based on		R package compcodeR
Number of repetitions	20	100 50
Samples per group	1 {2,5,8}	3
Total number of genes	16,500 28,800	10,000 10,000
UE genes	(1000,500) (1000,500), (800,1500)(1000,800), (2000,1000)–	–
Properties of DE genes (excl. UE)	{0.3, 0.5, 0.7}{0.6 + {0.1, 0.3, 0.5}	{0.25, 0.05} {0.45,0.55, 0.65,0.75}
Log ₂ fold-change	2 2 + 3	2 ≥ 3 2
Properties of upregulated genes	0.9 0.9 + 0.1	1 {0.5, 0.7, 0.9, 1}{0.5, 0.7, 0.9, 1}
Heterogeneous data	No Yes	No
Prefiltering	Mean count ≤ 2	Total count < 1
Competing methods	Type and number	3: edgeR, DESeq2, TCC
Standard edgeR	5: edgeR, DESeq/DESeq2, HTN, Library Size, NOISeq	Yes
Other method parameters	No (binomial test) No (based on Poisson)	See the original paper
Evaluation criteria	Metric	AUC (smoothed ROC)AUC
Aggregation	Boxplots	Boxplots

Note: Included are only components (i) for which the corresponding results are presented as figures or tables in the main paper (i.e., not in the supplement), (ii) that compare the method of interest to other competing methods, and (iii) that correspond to the performance evaluation based on simulated data. In addition, some components are not included in the experiment, which are indicated by asterisks (*). Competing methods and evaluation criteria for datasets not included in the experiment are not shown. In the case of design-implementation gaps, the table refers to the code for reproducing the results.

two methods, O21 also considers the less well-known method TCC (tag count comparison; Sun et al., 2013), arguing that it is not sufficient to compare a newly proposed method to the most commonly used methods (edgeR and DESeq2) as those might not be the ones best suited for the analysis. Moreover, they see TCC as the main alternative to their proposed method since the normalization algorithm used by TCC corresponds to the normalization algorithm used by one version of MBCdeg.

Interestingly, Z21 and O21 use different implementations of edgeR. While the implementation by O21 corresponds to one of the edgeR standard workflows, Z21 uses three different implementations of edgeR across their simulation settings of which only one would be typically considered as edgeR (but still with different parameters than O21), while the other two are only edgeR-like. One reason for this choice is that some simulation settings in Z21 do not have biological replicates (i.e., $n_{obs} = 1$ in each group), for which the standard edgeR implementation yields an error (see Supporting Information Section B.2 for details). Regarding the implementation of DESeq/DESeq2, Z21 actually use both DESeq and DESeq2, although they generally refer to the method as DESeq, the predecessor method of DESeq2. This might be explained by the fact that, similar to edgeR, DESeq2 is not intended for settings without biological replicates and thus yields an error, which is why Z21 uses DESeq in these settings. Note that it has been shown that DESeq and DESeq2 perform differently (Love et al., 2014). Both Z21 and O21 use the same parameters for DESeq2. For the parameters of the remaining methods see Z21 and O21 as well as the referenced code.

Evaluation criteria Both Z21 and O21 assess the methods' ability to correctly identify DE genes using the area under the receiver operating characteristic curve (AUC). They both justify this decision with the fact that the AUC, in contrast to other popular measures, does not require the choice of a threshold value. The AUC takes values from 0 to 1, where 1 corresponds to perfect discrimination of DE and non-DE (i.e., nondifferentially expressed) genes, and 0.5 corresponds to random assignment. However, due to an unfortunate default option in the R package used by Z21 to calculate the AUC, the resulting AUC values are 1 minus the correct AUC for some repetitions (see Supporting Information Section B.3 for details). Apart from the different R packages used to calculate the AUC, Z21 also employs a *smoothed* ROC curve (receiver operating characteristic curve) to estimate the AUC in some of their simulation settings (study 5), which can lead to slightly different results compared to the nonsmoothed ROC curve. Regarding the aggregation of AUC values across the simulation repetitions, both Z21 and O21 use boxplots.

4.2 | Challenges when conducting the experiment

Reproducibility When reproducing the results presented in O21 and Z21, we do not modify the original code in a way that would change the results, with one exception: We change the number of simulation repetitions from 10 to 20 (i.e., the number reported in the paper) in the code provided by Z21 since the results using 20 repetitions are more similar to the results shown in Z21 (note that we make this change before crossing the designs). As stated in Section 2.3, we also use the same R and R package versions as in the original papers (see Table S6 in the Supporting Information). However, Z21 does not provide this information, which is why we use the most recent package versions available when conducting the experiment (see our code documentation for the exact version information). The code by Z21 also does not include a random number seed, which we therefore set but which is most likely different from the seed used by Z21. Note that for reproducing the results of Z21, we use their AUC implementation potentially yielding incorrect results, but additionally calculate the correct version.

Based on these modifications, running the code of Z21 and O21 results in very similar but not exactly the same boxplots as shown in the original papers. More specifically, the relative performance of each method is the same in the original and reproduced versions, but some boxplots have, for example, different outliers. For Z21, this relatively high degree of reproducibility is noteworthy considering the fact that the provided code does not include a seed or version information. The only three settings that do not yield similar results are the settings from study 5 by Z21 (the differences between the original and reproduced results are described in Section 4.3). Apart from the aforementioned missing seed and version information, the different results in these settings could be due to the fact that the code might not have been provided in its final version.

Crossing the designs As already stated in the first example on cancer subtyping, conducting the cross-design experiment implies that all considered methods are applied to new datasets (new in the sense that these datasets have not been included in the original paper). It is thus necessary to carefully specify the method parameters of SFMEB, MBCdeg, and all competing methods. Although the simulation settings of Z21 and O21 are less comparable than the real datasets of N19 and R19 in the cancer subtyping example, we nevertheless adopt the parameter values from the original papers because we

consider the risk of running the methods with suboptimal parameter settings to be lower for the parameters used by Z21 and O21 than for parameters selected by ourselves (especially because we select the parameters *before* seeing the results to avoid the risk of favoring one of the methods, as stated in Section 2.3). However, both Z21 and O21 consider more than one parameter value for some methods, and Z21 even uses different methods across the simulation settings (i.e., DESeq and DESeq2). For all methods evaluated in Z21 (i.e., SFMEB and its competing methods), we adopt the parameters from study 5 since they are the most similar to the simulation settings considered in O21 (i.e., nonheterogeneous data, generated using the binomial distribution, with replicates). In all simulation settings of O21 included in our experiment, the authors evaluate two versions of MBCdeg, which are denoted as MBCdeg1 and MBCdeg2 and correspond to two different normalization options. Since MBCdeg1 and MBCdeg2 are also implemented separately in the code, we include both versions in the experiment but decide to focus on MBCdeg2, which was observed to be slightly more stable and accurate in O21, before seeing any results. Although O21 does not vary any other parameters of MBCdeg or the competing methods, we note that the main parameter of MBCdeg that is extensively discussed by O21 might not be ideal for some simulation settings of Z21. We thus conduct a sensitivity analysis using two different values for this parameter (see Section B.4 for details).

Since O21 and Z21 use the same evaluation criterion (i.e., boxplots representing the AUC values of all simulation repetitions), we only reevaluate the performance of SFMEB and MBCdeg on each other's competing methods and data. When crossing the designs, we do not consider the AUC that is based on the smoothed ROC curve used by Z21 in some simulation settings. Of course, we also do not use the incorrectly calculated version of the AUC.

Note that not all design components of Z21 and O21 are compatible. More specifically, the DESeq2 and edgeR implementation in O21 results in an error when applied to the simulation settings without biological replicates by Z21. As stated in Section 4.1, this is because DESeq2 and edgeR are not intended for settings without biological replicates and O21 does not use a (possibly nonideal) workaround solution as done by Z21.

4.3 | Results

Performance based on the original study design Figure 3A and 3D shows the reproduced performance results of SFMEB and MBCdeg2 with an additional dashed line corresponding to the median AUC of the corresponding method of interest overall simulation repetitions. Note that the method labels are adopted from the original papers although the competing methods DESeq and edgeR in Z21 do not exactly correspond to the actual method in some simulation settings as discussed above.

For SFMEB, we show both the reproduced AUC values that are potentially biased towards higher values and the correct AUC values. As stated in the previous section, we only observe a noteworthy performance difference between the reproduced results and the results shown in Z21 for three simulation settings (i.e., study 5). In these settings, two competing methods consistently show better performance in the reproduced version, leading to SFMEB being the second best instead of the best performing method in two settings. However, these differences become irrelevant when looking at the correct AUC results. In fact, only the AUC values of the competing methods are in some settings affected by the incorrect AUC calculation, resulting in SFMEB outperforming its competing methods more clearly than initially claimed by its authors. The performance results of SFMEB based on the corrected AUC values are thus still consistent with the conclusion of Z21 that SFMEB outperforms its competitors in most settings (achieving rank 1 according to median AUC in 13 out of 15 settings).

MBCdeg2 also performs well in its original study design. As noted by O21, the method tends to achieve higher AUC values in the settings with a small (≤ 0.45) proportion of DE genes. In some settings where the proportion of DE genes is ≥ 0.55 ; however, the method seems to fail, often resulting in AUC values below 0.25 and not being able to outperform any of its competing methods (the same applies to MBCdeg1). O21 discusses the occasional failure of MBCdeg extensively and concludes that the identification of the non-DE gene cluster (which they state to be the key to the proposed framework) fails in these cases, which leads to an incorrect classification of DE and non-DE genes. However, MBCdeg2 generally performs better than the competing method TCC in settings where TCC performs well (the same applies to MBCdeg1). Given the fact that TCC could be expected to outperform other methods since the datasets are generated using the TCC R package and the normalization algorithm used by TCC was designed for settings with asymmetric (i.e., $\neq 0.5$) upregulation as considered by O21, O21 see this as the main contribution of their study.

Performance based on the crossed design Figure 3B and 3C displays the performance results of MBCdeg2 and SFMEB based on each other's simulation data and competing methods. In the study design of O21, SFMEB generally shows worse

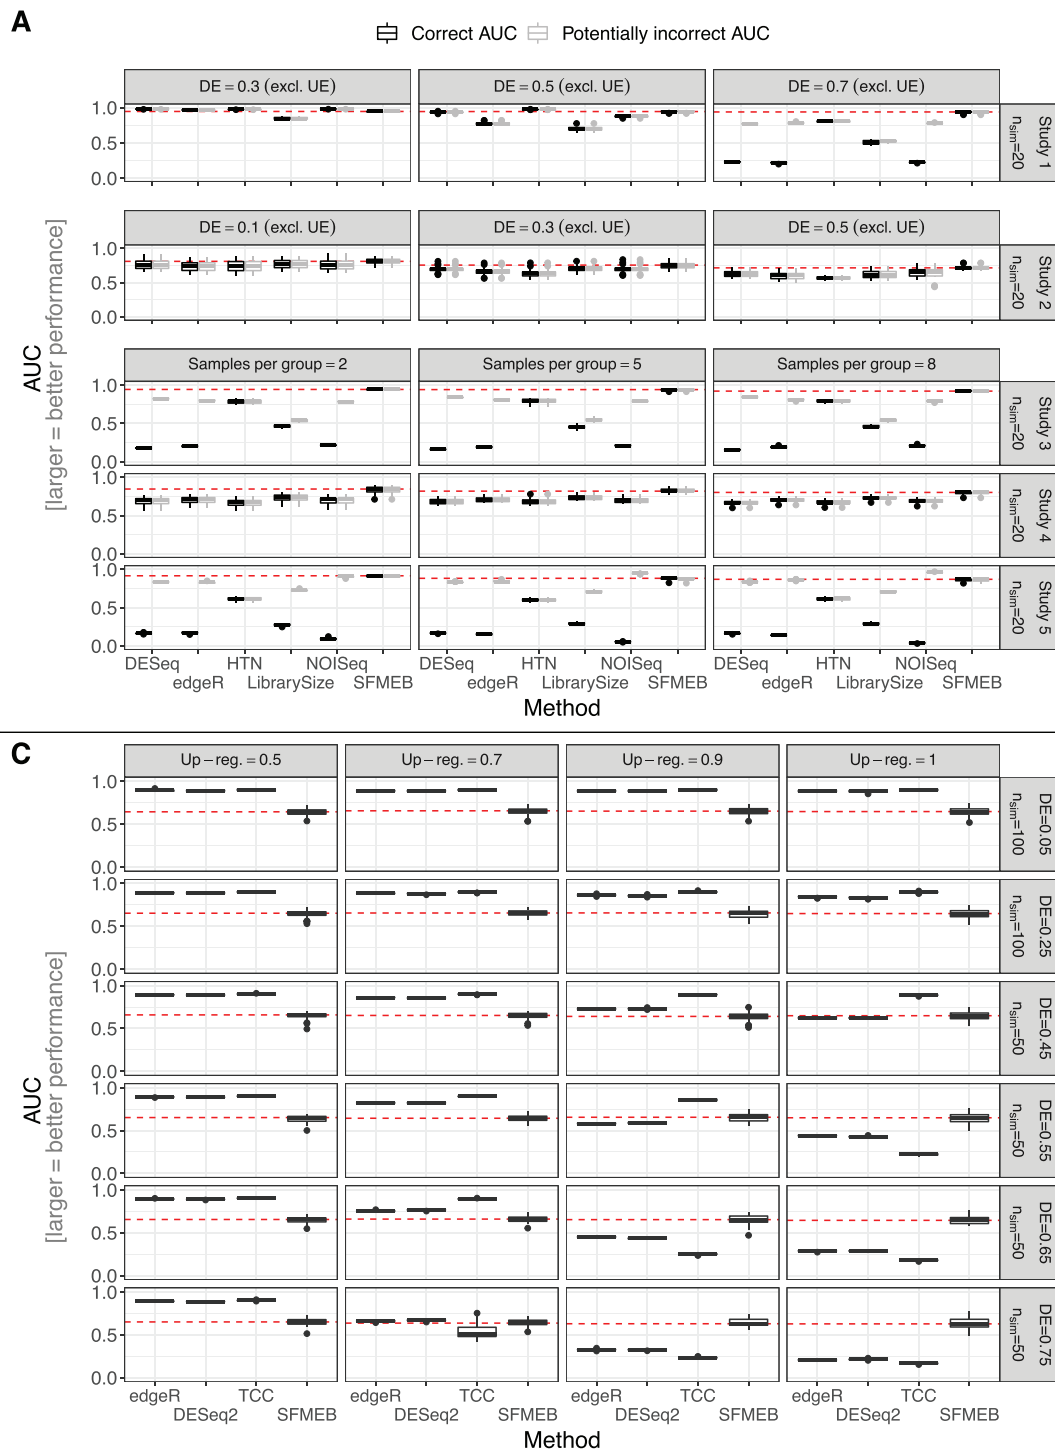


FIGURE 3 Results of the cross-design validation experiment for the differential gene expression analysis example, where the subfigures A–D correspond to the cells of Table 1. (A) Performance of SFMEB based on the design by Zhou et al. (2021) (= original design). (B) Performance of MBCdeg2 based on the design by Zhou et al. (2021) (= crossed design). (C) Performance of SFMEB based on the design by Osabe et al. (2021) (= crossed design). (D) Performance of MBCdeg2 based on the design by Osabe et al. (2021) (= original design). In each subfigure, the boxplots correspond to n_{sim} simulation repetitions, where $n_{sim} \in \{20, 50, 100\}$. The red dashed line corresponds to the median AUC of SFMEB (A and C) and MBCdeg2 (B and D) across all simulation repetitions. In the original paper by Zhou et al. (2021), the AUC has not been calculated as intended by the authors, which is why in subfigure A, both the correct AUC values and the reproduced and potentially incorrect AUC values are provided.

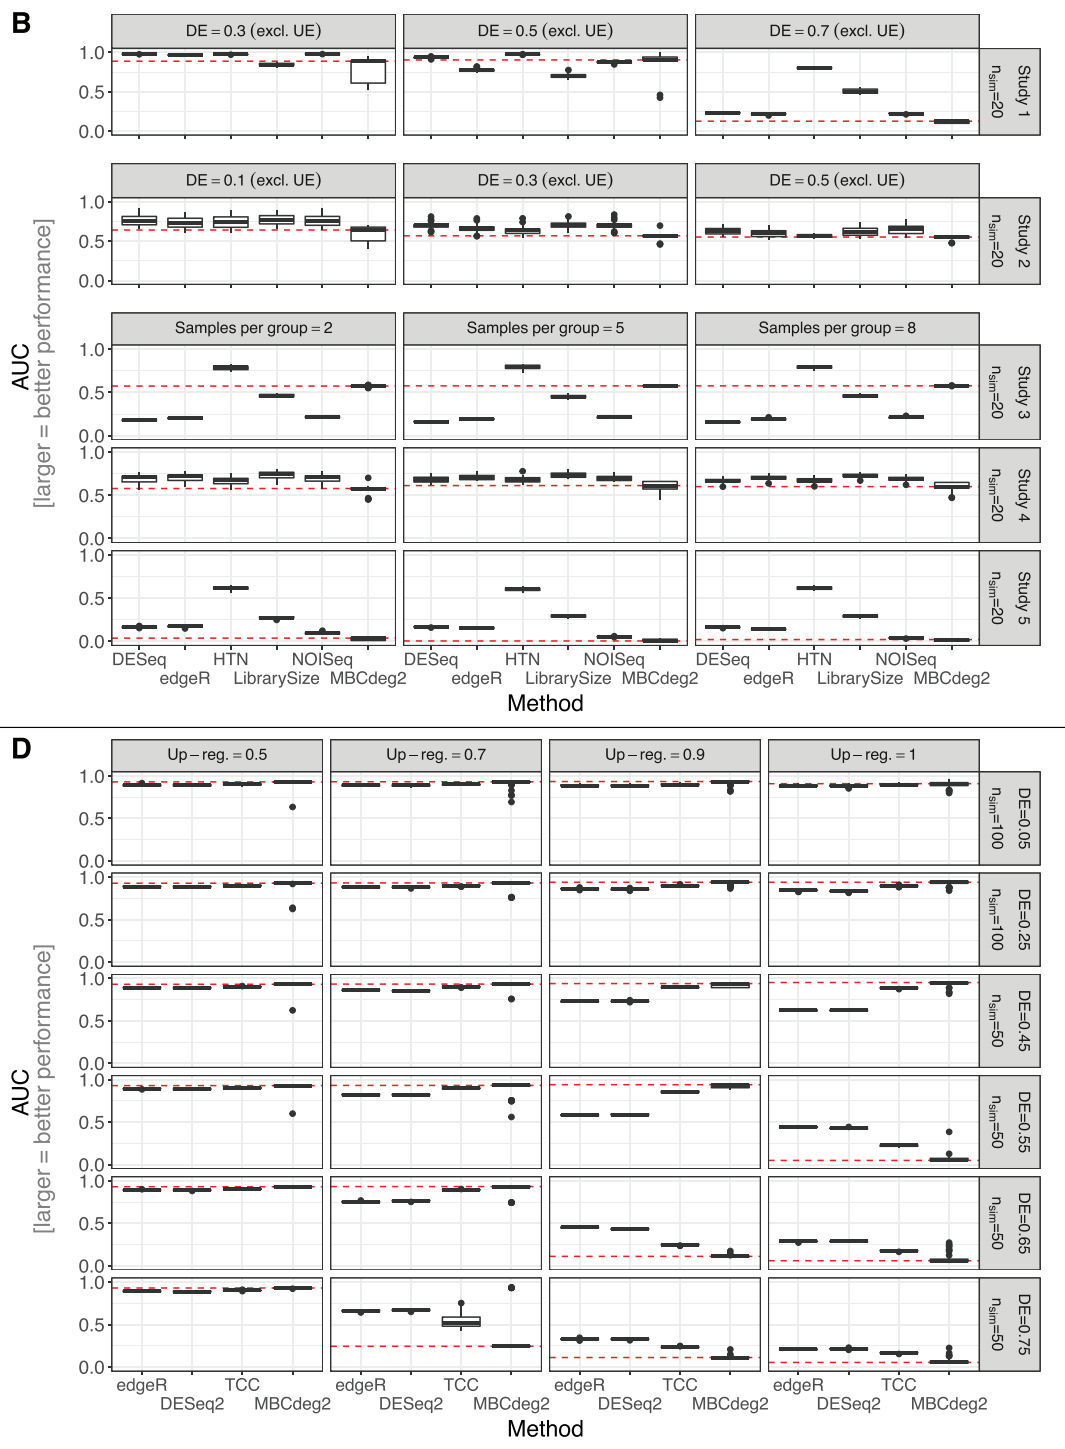


FIGURE 3 Continued

performance than in its original design, having lower median AUC values than all of its competitors in 17 out of 24 settings. However, in five out of the remaining seven settings (the settings with a high proportion of DE genes that are mostly upregulated in one group), SFMEB clearly outperforms the competing methods. Interestingly, this difference in relative performance is mainly caused by the varying AUC values of the competing methods edgeR, DESeq2, and TCC. SFMEB itself, on the other hand, shows very robust AUC values across all settings. However, with a median AUC of about 0.65 in each setting, SFMEB’s absolute performance is worse than in the original study, where the lowest median AUC of SFMEB is 0.72.

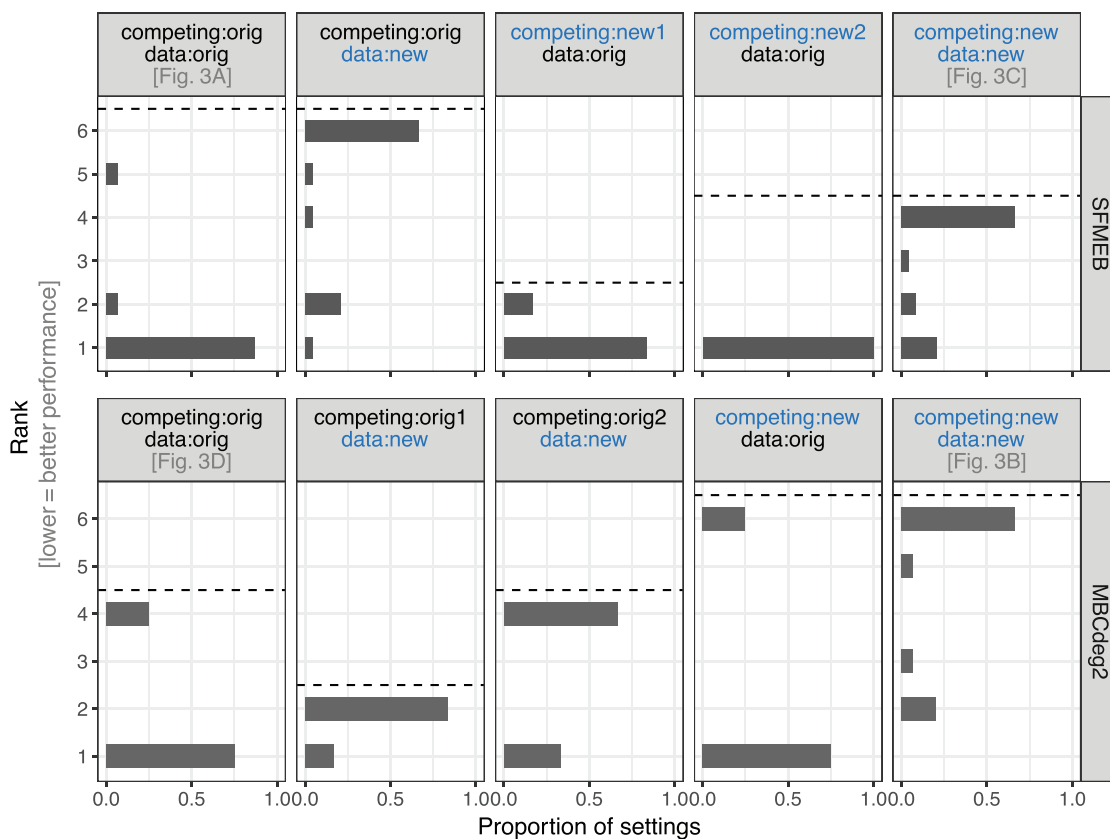


FIGURE 4 Performance ranks of the two differential gene expression analysis methods SFMEB and MBCdeg2 based on datasets and competing methods that either correspond to the original (SFMEB: Zhou et al., 2021; MBCdeg2: Osabe et al., 2021) or crossed design (SFMEB: Osabe et al., 2021; MBCdeg2: Zhou et al., 2021). Each horizontal bar plot shows the performance rank distribution of SFMEB or MBCdeg2 for one combination of datasets and competing methods. The number of ranks that is represented by each bar plot corresponds to the number of simulation settings in the respective data source ($n_{\text{setting}} = 15$ for data based on Z21; $n_{\text{setting}} = 24$ for data based on O21). The rank of each simulation setting is calculated based on the median AUC value across all simulation repetitions. The dashed lines indicate the number of compared methods, that is, the highest possible rank. Note that for both SFMEB and MBCdeg2, the ranks that result from the combination of competing methods by O21 and data by Z21 are represented by two bar plots due to the incompatibility of two competing methods of O21 with some simulation settings of Z21.

Similar to SFMEB, MBCdeg2 generally performs worse compared to its original design. In 10 out of 15 settings, it is outperformed by all competing methods. However, it is the second best method in four of the remaining five settings (based on median AUC). In contrast to SFMEB, the absolute performance varies more across the settings and only reaches a value comparable to the original design (excluding the settings where the method failed) in two settings. Similar to its original design, there are four settings where MBCdeg2 shows extremely low AUC values, which again seems to be caused by the incorrect identification of the non-DE cluster (note that these are all settings where the proportion of DE genes is ≥ 0.6 , which is consistent with O21's observation in the original paper). As stated in Section 4.2, we also conduct a sensitivity analysis where MBCdeg2's main parameter is set to a different value. However, this does not improve the AUC values (see Figure S3 in the Supporting Information).

Figure 4 shows the resulting performance ranks of SFMEB and MBCdeg2 when data and competing methods are varied individually. For both SFMEB and MBCdeg2, the datasets and competing methods can either be set to the original or the crossed version, which results in four ($= 2^2$) different study designs. Note that for both SFMEB and MBCdeg2, the study design that is based on the competing methods of O21 and the data of Z21 are represented by two panels instead of one. This is because two of the three competing methods of O21 cannot be run in some simulation settings of Z21 (see Section 4.2), which makes the resulting ranks incomparable to the ranks that are based on the simulation settings with all three competing methods. Within each study design, the ranks are calculated separately for each simulation setting based on the median AUC and are summarized as bar plots (i.e., each bar plot displays the distribution of 15 or 24 ranks, which corresponds to the total number of settings considered by Z21 and O21, respectively). All AUC values are calculated

correctly. For both SFMEB and MBCdeg2, the performance mainly depends on which simulated datasets are considered. In contrast, using different competing methods has no considerable impact on the distribution of ranks, except that the maximum possible rank reflecting the worst method varies according to the number of competing methods. This is also due to the partial overlap of competing methods between Z21 and O21.

The results of MBCdeg1 based on the crossed design are very similar to the results of MBCdeg2 and can be found in the Supporting Information file (Figure S4).

5 | DISCUSSION

5.1 | Summary of results and limitations

In this paper, we conducted a systematic experiment, which we refer to as “cross-validation of methods” and in which we reevaluated methods based on the datasets, competing methods, and evaluation criteria of a paper proposing a method for the same data analysis task. We considered two exemplary data analysis tasks, namely cancer subtyping using multiomic data and differential gene expression analysis. For each analysis task, we selected two methods, PINSPlus (Nguyen et al., 2019) and NEMO (Rappoport & Shamir, 2019) for cancer subtyping, and SFMEB (Zhou et al., 2021) and MBCdeg (Osabe et al., 2021) for differential expression analysis.

Although we did not conduct our cross-design validation experiment on a large scale, several interesting findings emerged. First, the difficulties in finding eligible papers showed that many papers are still being published without openly available code to reproduce the results. For the papers that were selected, running the provided code did not yield the exact same results as presented in the respective paper. Only the results of PINSPlus were close to being fully reproducible with only one differing p -value in one of the competing methods. Although the lack of reproducibility could be partly due to, for example, our computational resources that were different from those of the authors of the four papers, other potential reasons are that some codes were not provided in their final version and that not all R/R package versions were reported. The latter is particularly relevant for R, which is subject to frequent package updates (potentially causing errors or changing results) and, in contrast to the programming language Stata, does not have integrated version control. Nevertheless, the reproduced results of all four methods were consistent with the conclusion of the original papers that the respective method shows good performance.

Second, the experiment concretely illustrated the researchers’ degrees of freedom regarding the performance assessment of a method. Notably, all four study designs seemed well thought-out and the authors provided justifications in most cases. Interestingly, even for the design components that were similar in both papers, the exact implementation was often different. For example, SNF and edgeR were included as competing methods in both papers of the cancer subtyping and differential expression analysis task, respectively, but were run with different parameters.

Third, the experiment showed how differences in the study design can affect the performance of a method. Three out of the four considered methods (PINSPlus, SFMEB, and MBCdeg) performed worse when assessed on the crossed study design, which seems to be consistent with the general concern that the performance of newly proposed methods is overoptimistic (Boulesteix et al., 2013; Buchka et al., 2021; Norel et al., 2011). Only one method, NEMO, performed well when evaluated on the study design of PINSPlus’ original paper and only showed slightly worse performance in some settings where datasets, competing methods, and evaluation criteria were varied individually. For both analysis tasks, using different datasets (real or simulated) had the largest impact on the performance results, which was particularly surprising for the real datasets of the cancer subtyping example where both papers used the same data type and source.

It is important to note that while the findings of our experiment might help to see the performance reported in the original papers from a different perspective, they cannot be seen as evidence of any of the four methods generally having good or bad performance. First, our experiment is limited in the sense that we did not include all study designs and corresponding results reported in the papers, which gives an incomplete picture regarding the study design of the papers and, importantly, the individual strengths and weaknesses of each method. This also includes qualitative evaluation criteria such as PINSPlus’ user-friendliness regarding the choice of the number of clusters (which was also noted by Duan et al., 2021), NEMO’s simplicity and support of partial data, the avoidance of potential error-prone data normalization when using SFMEB, and the high interpretability of MBCdeg’s main parameter. Second, the method performances observed in the experiment clearly depend on (i) our own expertise regarding each method and (ii) the respective new design we reevaluated each method on. The latter is the result of an informal and unsystematic search process based on eligibility criteria (i.e., R as a programming language and publicly available code) that could have been specified differently. In

addition, reevaluating each method on more than one new design (i.e., extending the 2×2 table in Section 2 [Table 1] to a $K \times K$ table) could lead to different and more nuanced results.

In addition to the restricted informative value regarding the performance assessment of the considered methods, our experiment is also limited in the sense that the deteriorating method performance observed in three of the four methods cannot be transferred to methodological research in general. This is due to the fact that we only considered two data analysis tasks and, as mentioned above, only included two papers per data analysis task that were selected based on an informal search process.

5.2 | Mechanisms leading to an optimistic performance evaluation and possible solutions

Although the results of our experiment cannot be seen as general evidence for the optimistic performance evaluation of newly proposed methods in methodological research, the experiment itself provides insights into the mechanisms that might explain the observed performance differences. In the following, we will discuss four of these mechanisms, which have either been addressed frequently in the literature or are rarely mentioned in the literature but seem to have been present in our experiment. In addition, we point to possible solutions that can help to avoid large performance discrepancies between original and subsequent studies.

Overfitting of study design to method Our experiment illustrated the many degrees of freedom existing in the assessment of a method's performance. This flexibility can tempt researchers to choose the study design in favor of their proposed method. This may happen both at the planning stage when researchers primarily select a study design in which their method is expected to perform well (e.g., leaving competing methods at their default parameters or simulating data from the model underlying the proposed method), and after seeing the results when they add and/or omit certain design components (e.g., simulation parameters or evaluation criteria; Nießl et al., 2022; Pawel et al., 2022; Ullmann et al., 2023). Focusing on advantageous designs at the planning stage is not necessarily a questionable research practice but becomes problematic if not clearly stated. Changing the study design *after* seeing the results may be legitimate in some cases as far as it is transparently reported, for example, if the originally chosen evaluation criterion turns out to behave inadequately for all methods. But changing the study design *is* bad practice if it is performed in a cherry-picking fashion, that is, excluding or including results depending on whether they convey the expected message or not. The “overfitting” of the study design to the method increases the risk of obtaining different, less optimistic conclusions in a subsequent comparison study in which the authors have less incentives to present the corresponding method in a favorable light.

As already noted by Simmons et al. (2011) in the context of applied research, such optimizations most often do not reflect malicious intent. Instead, they are usually the result of self-serving interpretations of ambiguity convincing honest researchers that the decisions (in our case, regarding the study design) matching their expectations and hopes are the most appropriate ones for various other reasons. These mechanisms are certainly encouraged by publication pressure and publication bias (Boulesteix et al., 2017). Selective reporting after seeing the results can be largely avoided by preregistering study designs and documenting all changes that have to be made subsequently (Morris et al., 2019; Pawel et al., 2022). However, it does not prevent authors from selecting advantageous designs from the start when planning their study. This pitfall could be avoided by adapting the designs from previous studies conducted by different authors. Although designs from different studies might not be suitable to demonstrate all features of the new method, the inclusion of at least one setting that is more “fair” for all compared methods and does not obviously favor the new method (even if the setting is generated by the authors themselves and not adapted from a different study) reduces the risk of overoptimistic conclusions.

For the papers considered in our study, we do not assume that any components regarding the datasets, competing methods or evaluation criteria have been optimized to make the corresponding method of interest appear better than it actually is. On the other hand, we cannot completely rule out this possibility, although it is especially unlikely for NEMO, which was evaluated using a study design adopted from a previously conducted comparison study (Rappoport & Shamir, 2018), similar to preregistration where the design is fixed in advance.

Overfitting of the method to study design Just as the study design can be “overfitted” to the method of interest, the method of interest can also be “overfitted,” that is, overoptimized to the study design. This was already noted by Jelizarow et al. (2010) and Ullmann et al. (2023) with a focus on overfitting to the considered datasets. Since method development is, in itself, an optimization process that usually consists of several improvements after seeing the performance results, it is difficult to determine the point where further optimization amounts to overfitting the method to the design used for performance assessment. This not only concerns the method characteristics that are not intended to be changed by the

user but also the parameters that can be set by the method user and whose optimal values for different settings might also be overfitted to the considered study design (Pawel et al., 2022; Ullmann et al., 2023). Note that the issue of overfitting of the method to the study design is relevant for any method evaluation whose results are to be generalized to other evaluation criteria or datasets. This also includes methods that are developed for very specific applications, as long as the authors of the method want it to be used for any other evaluation criteria or datasets than those used for performance assessment (at least for the datasets, this usually seems to be the case).

To avoid overfitting of the method of interest to the study design, it is recommended to evaluate the method extensively. This includes using a large number of datasets and/or simulation settings and several evaluation criteria as well as checking the robustness of the method with respect to small changes in the study design since this makes it more difficult for the method to be artificially optimized (Boulesteix, 2015; Nießl et al., 2022; Norel et al., 2011; Ullmann et al., 2023). In principle, this is comparable to the classical context of regression where overfitting is less likely to occur if the number of observations is large.

Moreover, it may be helpful to reevaluate newly developed methods using a different design after the termination of the trial-and-error process, which might yield slightly worse but likely more realistic performance results (in the sense that the performance discrepancy between original and subsequent papers decreases). Although previous literature usually focuses on evaluating the method on new data (Jelizarow et al., 2010; Norel et al., 2011; Ullmann et al., 2023), considering different competing methods and evaluation criteria could also be reasonable. To reduce the risk of choosing the new design in favor of the proposed method, one could apply the design of a previous study conducted by different authors. As discussed above, the design of a previous study might not be suited to present all features of the proposed method (or even fully match the method's potentially very specific field of application) but this might be less relevant if the design is considered as an additional "external validation design". An external validation design could be, for example, the design of a neutral comparison study, or, similar to our experiment, a previously proposed method (e.g., a method that was included as a competing method). This procedure is only feasible without much additional effort if the authors of the previous paper have made the code for reproducing the results openly available and does not protect against systematic manipulation (e.g., modifying the method after seeing the results and thus consciously biasing the external validation).

When reading a paper, it is typically not possible to identify whether the method of interest has been overfitted to the design used for method development and, unless explicitly stated, if there are any settings that have been separated from the development process. This also applies to the papers included in our experiment, which do not have a corresponding statement. However, MBCdeg is mainly based on an algorithm that was developed by different authors for a different analysis task (i.e., clustering of genes that have already been identified as differentially expressed), which means that this part of the method cannot be overfitted to the design of Osabe et al. (2021).

Different levels of expertise While the mechanisms discussed above are mostly attributed to the nonneutrality of the authors proposing their new method, there are also other potential mechanisms leading to deteriorating performances in subsequent papers. One of them originates from the fact that, as already noted by Duin (1996), the performance of a method is not just dependent on the design it is evaluated on but also on the skill of the person who applies the method. The difference in performance between original and subsequent papers may therefore also be due to the lower expertise level of the subsequent authors whose parameter choice when applying the method to the new data is likely to be less optimal than the parameters that the authors of the method would have selected. Of course, the degree to which the performance deteriorates due to the lack of expertise may be different for each method (Boulesteix et al., 2017) and also depends on how much the new design in which the method is applied differs from the design of the original paper.

As described in Sections 3.2 and 4.2, we also faced the challenge of choosing appropriate method parameters when applying the methods of our experiment to the new datasets and we cannot rule out that these decisions might have led to a worse performance than if the authors of the original papers had chosen the parameters themselves. In the first example on cancer subtyping, we note that although the datasets in both papers had the same data type and originated from the same source, the authors of NEMO and PINSPlus might have set different parameters (including method specific preprocessing steps) for their respective method since the datasets have a different distribution of samples and omic variables (due to the different preprocessing steps and number of datasets). For example, the authors of PINSPlus might have normalized the data when applying it to the datasets of NEMO. The same applies to the differential expression analysis example, where we decided to set SFMEB's parameters for the crossed simulation data as in the simulation setting of the original paper that seemed to be the most similar to the new simulation. It is possible that the authors of SFMEB who are experts in this method might have used a different parameter setting. For MBCdeg, we also cannot rule out that our low level of expertise has contributed to the deteriorating performance of the method. Although we evaluated different values for one parameter of MBCdeg as a sensitivity analysis, we only did that to a limited extent and the considered values may still

be suboptimal (e.g., the authors did not specify how the parameter should be set in the presence of uniquely expressed genes, which are not considered in their simulation settings). It also has to be noted that we are nonexpert users for many of the competing methods used for each paper, and, for instance, the performance of Consensus Clustering and iCluster+ (competing methods of PINSPPlus) is certainly dependent on the expertise level of the user since the optimal number of clusters has to be specified manually based on different types of plots and is thus very subjective. However, the difference in expertise (i.e., comparing our expertise vs. the expertise of the authors of the four papers) is probably less drastic with regard to the competing methods than for the methods of interest and is thus not of equal relevance.

One possibility to avoid the systematic deterioration of performance in subsequent studies due to a lower level of expertise is to involve the authors of the method in the respective study (Boulesteix et al., 2017; Morris et al., 2019; Pawel et al., 2022). This can be realized if they implement their method themselves, as done, for example, in the study by Zapf et al. (2021) that involved the authors of all considered methods as co-authors or in benchmark studies that are organized as challenges such as the DREAM challenges (<https://dreamchallenges.org/>). Alternatively, the authors of a method can be contacted to make sure that their method is implemented correctly as done in the comparison study by Herrmann et al. (2021). However, while the authors of a method could potentially be involved in the majority of comparison studies that assess their method, they will not be able to verify the correct implementation of their method in every *applied* study. Although there is value in studying the performance of a method when used by an expert, it might thus be even more important to assess the performance when it is applied by nonexperts (Boulesteix et al., 2017; Duin, 1996), as we did in this experiment. Note, however, that even among the nonexperts of a method, there are different levels of expertise—or a different willingness to gain expertise by getting more familiar with the method (which may apply in particular to authors that use the method as a competitor for their own method).

In general, it might thus be advisable for authors to make the performance of their method less dependent on user expertise by providing high-quality method documentation that includes a description of all method components and parameters, concrete guidelines on how to choose optimal parameter values in different applications, and ideally also tutorials that help users to become more familiar with the method (Bokulich et al., 2020). If feasible, method authors can also implement automated parameter selection, which protects against the above-mentioned tendency to leave method parameters of competing methods at default values. Moreover, reporting the robustness of the method performance with respect to different parameter values (as done by all four papers considered in the experiment) allows method users to gain an understanding of which parameters need to be carefully specified (Ullmann et al., 2023). Of course, reducing the effect of different levels of expertise also requires efforts from the authors of subsequent papers who need to consider the available guidelines and information on how to set the method parameters.

Different fields of application An insight we gained from the experiment that seems to be rarely addressed in the literature but plays an important role in the optimistic performance evaluation of newly proposed methods is related to the appropriate field of application of a method and its individual strengths within this field. If a method performs worse in a subsequent paper, this can indeed be due to the mutual overfitting of method and design or the lack of expertise, as discussed above. However, the deteriorating performance could also be explained by the fact that the field of application of the subsequent study does not exactly match the field of application the method is intended for. Unfortunately, our experiment suggests that it is often hard to assess if this is the case.

For example, although NEMO and PINSPPlus obviously have the same *general* field of application (i.e., cancer subtyping using multiomic data), it was clear that PINSPPlus, in contrast to NEMO, is not intended to be used on partial multiomic datasets (i.e., datasets where some patients do not have any measurements for one or more omic data type), which is why we excluded them from our experiment. On the other hand, PINSPPlus was initially (i.e., in its original paper) only evaluated based on its ability of finding subtypes that have significantly different survival while NEMO was additionally assessed based on the enrichment of certain clinical variables such as the tumor stage. We did not exclude the clinical enrichment criterion, although it could be argued that PINSPPlus is only intended for applications where it is relevant to find subtypes with different survival. Similarly, in the differential expression analysis example, we excluded the three-group simulated data used to assess the performance of MBCdeg in the original paper since the authors of SFMEB did not explicitly mention that their method is intended for this type of application. On the other hand, we did not exclude the settings without biological replicates (i.e., $n_{obs} = 1$ in each group) used by the authors of SFMEB from our experiment although the authors of MBCdeg did not explicitly state that settings without biological replicates belong to MBCdeg's field of application (and other popular methods such as edgeR and DESeq2 are explicitly not intended for these settings). Moreover, it is not clear whether MBCdeg can be applied in settings with uniquely expressed genes (i.e., genes with zero counts in one condition), which were included in most settings used to evaluate SFMEB.

These examples show that it is often not clear for method users what the method's exact field of application is, which consequently makes decisions on whether it is appropriate to apply the method to a new design more difficult and subjective. On the other hand, authors proposing a new method cannot be expected to provide an exact definition of the method's appropriate field of application that accounts for every imaginable design, and some authors explicitly state that the method simply requires more evaluation in certain designs to assess whether they belong to the method's appropriate field of application. For example, the authors of MBCdeg mention that their method still needs to be evaluated on additional simulation frameworks and real data with different experimental settings and organisms.

In general, authors proposing a new method should thus try to study and report its field of application as comprehensively as possible, which, in addition to guidelines for choosing adequate method parameters discussed above, we also consider an important part of the method documentation. As a means to this end, authors should investigate their method's performance in relation to the properties of the included datasets instead of focusing on its overall performance (Strobl & Leisch, 2022). On the other hand, authors using the method in a subsequent study should carefully check whether the application of the method is appropriate and ideally point to differences in the study design.

An issue related to the field of application is that methods often have specific strengths or features *within* their field of application, which is typically reflected by the design and not problematic if reported transparently (as discussed above). However, the method's strengths and special features may not be highlighted to the same extent through the design of the subsequent study (which may be, for instance, selected to highlight the strengths of a different method), thus leading to a deteriorating performance.

We also observed this in our experiment. As mentioned above, a special feature of NEMO is that it can handle missing values in the omic data. However, this feature does not come into play in the study design of PINSPlus, which cannot handle missing values (so that its authors did not consider designs with missing data). Notably, NEMO outperformed the competing methods in the original paper even more clearly for the datasets with missing values than for the full datasets, and although NEMO showed good performance in the design of PINSPlus, its performance might have been even better if the crossed design had also included datasets with missing values. In the differential expression example, the authors of SFMEB emphasize its strength of not requiring data normalization, which is an essential step for most other methods that can mislead downstream analysis if not done correctly. The authors of SFMEB include several data settings where normalization can be error-prone, such as heterogeneous datasets with clearly different fold changes between the conditions. This special strength is, however, not relevant for the settings of MBCdeg that are included in our experiment, which may have also led to SFMEB's deteriorating performance.

In contrast to the mismatch regarding the appropriate field of application discussed above, it is not necessarily inappropriate if a subsequent study disregards the strengths of a method, but it should be ideally mentioned. Note that the discussed mechanisms can also be applied to the competing methods of the original and subsequent papers, whose field of application and specific strengths might be more or less reflected by the study design.

6 | CONCLUSION

Based on the insights gained from the cross-design validation experiment, we conclude that while the discrepancy between original and subsequent studies assessing the performance of a method may be, in part, attributed to the nonneutrality of the method's authors, there are also other mechanisms related to different levels of expertise and fields of application that can contribute to a deteriorating method performance. It is important that both the authors proposing a method and the authors applying the method in a subsequent study acknowledge and counteract these mechanisms. On the side of the method authors, this requires not only a transparent and extensive evaluation but also comprehensive method documentation that enables correct usage by other researchers. In terms of transparency, a minimum requirement for all papers proposing and/or comparing methods should be to openly provide the code, software versions, computational environment, and, if possible, data to reproduce the results. This does not guarantee but at least facilitates the detection of potential overoptimistic statements in the original papers and the nonappropriate use of the methods in subsequent papers. In the long run, these efforts will increase the reliability of studies proposing new methods.

ACKNOWLEDGMENTS

This work was supported by the German Federal Ministry of Education and Research (01IS18036A) and by the German Research Foundation (BO3139/4-3, BO3139/7-1) to ALB. The authors of this work take full responsibility for its content. The authors thank Milena Wünsch for helpful comments and Anna Jacob for language correction.

Open access funding enabled and organized by Projekt DEAL.


CONFLICT OF INTEREST STATEMENT

The authors have declared no conflict of interest.

DATA AVAILABILITY STATEMENT

The R code and data to reproduce all results are openly available at <https://doi.org/10.6084/m9.figshare.20754028>. The R code without data is also available as supplementary material.

OPEN RESEARCH BADGES

 This article has earned an Open Data badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. The data is available in the [Supporting Information](#) section.

This article has earned an open data badge “**Reproducible Research**” for making publicly available the code necessary to reproduce the reported results. The results reported in this article could fully be reproduced.

ORCID

Christina Nießl  <https://orcid.org/0000-0003-2425-7858>

Anne-Laure Boulesteix  <https://orcid.org/0000-0002-2729-0947>

REFERENCES

- Anders, S., & Huber, W. (2010). Differential expression analysis for sequence count data. *Genome Biology*, *11*, 106.
- Bokulich, N. A., Ziemski, M., Robeson, M. S., & Kaehler, B. D. (2020). Measuring the microbiome: Best practices for developing and benchmarking microbiomics methods. *Computational and Structural Biotechnology Journal*, *18*, 4048–4062.
- Boulesteix, A.-L. (2015). Ten simple rules for reducing overoptimistic reporting in methodological computational research. *PLoS Computational Biology*, *11*, e1004191.
- Boulesteix, A.-L., Lauer, S., & Eugster, M. J. (2013). A plea for neutral comparison studies in computational sciences. *PLoS ONE*, *8*, e61562.
- Boulesteix, A.-L., Wilson, R., & Hapfelmeier, A. (2017). Towards evidence-based computational statistics: Lessons from clinical research on the role and design of real-data benchmark studies. *BMC Medical Research Methodology*, *17*, 138.
- Buchka, S., Hapfelmeier, A., Gardner, P. P., Wilson, R., & Boulesteix, A.-L. (2021). On the optimistic performance evaluation of newly introduced bioinformatic methods. *Genome Biology*, *22*, 152.
- Duan, R., Gao, L., Gao, Y., Hu, Y., Xu, H., Huang, M., Song, K., Wang, H., Dong, Y., Jiang, C., Zhang, C., & Jia, S. (2021). Evaluation and comparison of multi-omics data integration methods for cancer subtyping. *PLoS Computational Biology*, *17*, e1009224.
- Duin, R. P. W. (1996). A note on comparing classifiers. *Pattern Recognition Letters*, *17*, 529–536.
- Herrmann, M., Probst, P., Hornung, R., Jurinovic, V., & Boulesteix, A.-L. (2021). Large-scale benchmark study of survival prediction methods using multi-omics data. *Briefings in Bioinformatics*, *22*, bbaa167.
- Hofner, B., Schmid, M., & Edler, L. (2016). Reproducible research in statistics: A review and guidelines for the Biometrical Journal. *Biometrical Journal*, *58*, 416–427.
- Jelizarow, M., Guillemot, V., Tenenhaus, A., Strimmer, K., & Boulesteix, A.-L. (2010). Over-optimism in bioinformatics: An illustration. *Bioinformatics*, *26*, 1990–1998.
- Klau, S., Martin-Magniette, M.-L., Boulesteix, A.-L., & Hoffmann, S. (2020). Sampling uncertainty versus method uncertainty: A general framework with applications to omics biomarker selection. *Biometrical Journal*, *62*, 670–687.
- Lohmann, A., Astivia, O. L., Morris, T. P., & Groenwold, R. H. (2022). It's time! Ten reasons to start replicating simulation studies. *Frontiers in Epidemiology*, *2*, 973470.
- Love, M. I., Huber, W., & Anders, S. (2014). Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biology*, *15*, 550.
- Monti, S., Tamayo, P., Mesirov, J., & Golub, T. (2003). Consensus clustering: A resampling-based method for class discovery and visualization of gene expression microarray data. *Machine Learning*, *52*, 91–118.
- Morris, T. P., White, I. R., & Crowther, M. J. (2019). Using simulation studies to evaluate statistical methods. *Statistics in Medicine*, *38*, 2074–2102.
- Nguyen, H., Shrestha, S., Draghici, S., & Nguyen, T. (2019). PINSPlus: A tool for tumor subtype discovery in integrated genomic data. *Bioinformatics*, *35*, 2843–2846.
- Nguyen, T., Tagett, R., Diaz, D., & Draghici, S. (2017). A novel approach for data integration and disease subtyping. *Genome Research*, *27*, 2025–2039.
- Nießl, C., Herrmann, M., Wiedemann, C., Casalicchio, G., & Boulesteix, A.-L. (2022). Over-optimism in benchmark studies and the multiplicity of design and analysis options when interpreting their results. *WIREs Data Mining and Knowledge Discovery*, *12*, e1441.
- Norel, R., Rice, J. J., & Stolovitzky, G. (2011). The self-assessment trap: Can we all be better than average? *Molecular Systems Biology*, *7*, 537.

- Osabe, T., Shimizu, K., & Kadota, K. (2021). Differential expression analysis using a model-based gene clustering algorithm for RNA-seq data. *BMC Bioinformatics*, 22, 511.
- Pawel, S., Kook, L., & Reeve, K. (2022). *Pitfalls and potentials in simulation studies: Questionable research practices in comparative simulation studies allow for spurious claims of superiority of any method*. <https://doi.org/10.48550/arXiv.2203.13076>
- Rappoport, N., & Shamir, R. (2018). Multi-omic and multi-view clustering algorithms: Review and cancer benchmark. *Nucleic Acids Research*, 46, 10546–10562.
- Rappoport, N., & Shamir, R. (2019). NEMO: Cancer subtyping by integration of partial multi-omic data. *Bioinformatics*, 35, 3348–3356.
- Rigaill, G., Balzergue, S., Brunaud, V., Blondet, E., Rau, A., Rogier, O., Caius, J., Maugis-Rabusseau, C., Soubigou-Taconnat, L., Aubourg, S., Lurin, C., Martin-Magniette, M.-L., & Delannoy, E. (2018). Synthetic data sets for the identification of key ingredients for RNA-seq differential analysis. *Briefings in Bioinformatics*, 19, 65–76.
- Robinson, M. D., McCarthy, D. J., & Smyth, G. K. (2009). edgeR: A bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26, 139–140.
- Robinson, M. D., & Oshlack, A. (2010). A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biology*, 11, R25.
- Seyednasrollah, F., Laiho, A., & Elo, L. L. (2013). Comparison of software packages for detecting differential expression in RNA-seq studies. *Briefings in Bioinformatics*, 16, 59–70.
- Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E. C., Bahník, Š., Bai, F., Bannard, C., Bonnier, E., Carlsson, R., Cheung, F., Christensen, G., Clay, R., Craig, M. A., Dalla Rosa, A., Dam, L., Evans, M. H., Flores Cervantes, I., ... Nosek, B. A. (2018). Many analysts, one data set: Making transparent how variations in analytic choices affect results. *Advances in Methods and Practices in Psychological Science*, 1, 337–356.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359–1366.
- Sonabend, R., Bender, A., & Vollmer, S. (2022). Avoiding C-hacking when evaluating survival distribution predictions with discrimination measures. *Bioinformatics*, 38, 4178–4184.
- Soneson, C. (2014). compcodeR - an R package for benchmarking differential expression methods for RNA-seq data. *Bioinformatics*, 30, 2517–2518.
- Soneson, C., & Delorenzi, M. (2013). A comparison of methods for differential expression analysis of RNA-seq data. *BMC Bioinformatics*, 14, 91.
- Strobl, C., & Leisch, F. (2022). Against the “one method fits all data sets” philosophy for comparison studies in methodological research. *Biometrical Journal*. Advanced online publication. <https://doi.org/10.1002/bimj.202200104>
- Subramanian, I., Verma, S., Kumar, S., Jere, A., & Anamika, K. (2020). Multi-omics data integration, interpretation, and its application. *Bioinformatics and Biology Insights*, 14, 7–9.
- Sun, J., Nishiyama, T., Shimizu, K., & Kadota, K. (2013). TCC: An R package for comparing tag count data with robust normalization strategies. *BMC Bioinformatics*, 14, 219.
- Tepeli, Y. I., Ünal, A. B., Akdemir, F. M., & Tastan, O. (2020). PAMOGK: A pathway graph kernel-based multiomics approach for patient clustering. *Bioinformatics*, 36, 5237–5246.
- Ullmann, T., Beer, A., Hünemörder, M., Seidl, T., & Boulesteix, A.-L. (2023). Over-optimistic evaluation and reporting of novel cluster algorithms: an illustrative study. *Advances in Data Analysis and Classification*, 17, 211–238.
- Wang, B., Mezlini, A. M., Demir, F., Fiume, M., Tu, Z., Brudno, M., Haibe-Kains, B., & Goldenberg, A. (2014). Similarity network fusion for aggregating data types on a genomic scale. *Nature Methods*, 11, 333–337.
- Zapf, A., Albert, C., Frömke, C., Haase, M., Hoyer, A., Jones, H. E., & Rucker, G. (2021). Meta-analysis of diagnostic accuracy studies with multiple thresholds: Comparison of different approaches. *Biometrical Journal*, 63, 699–711.
- Zhou, Y., Yang, B., Wang, J., Zhu, J., & Tian, G. (2021). A scaling-free minimum enclosing ball method to detect differentially expressed genes for RNA-seq data. *BMC Genomics*, 22, 479.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Nießl, C., Hoffmann, S., Ullmann, T., & Boulesteix, A.-L. (2024). Explaining the optimistic performance evaluation of newly proposed methods: A cross-design validation experiment. *Biometrical Journal*, 66, 2200238. <https://doi.org/10.1002/bimj.202200238>