

Computational quality control tools for mass spectrometry proteomics

Wout Bittremieux^{1,2}, Dirk Valkenburg^{3,4,5}, Lennart Martens^{6,7,8}, Kris Laukens^{1,2}

¹Department of Mathematics and Computer Science, University of Antwerp, 2020 Antwerp, Belgium; ²Biomedical Informatics Research Center Antwerp (biomina), University of Antwerp/Antwerp University Hospital, 2650 Edegem, Belgium; ³Flemish Institute for Technological Research (VITO), 2400 Mol, Belgium; ⁴CFP, University of Antwerp, 2020 Antwerp, Belgium; ⁵I-BioStat, Hasselt University, 3590 Diepenbeek, Belgium; ⁶Medical Biotechnology Center, VIB, 9000 Ghent, Belgium; ⁷Department of Biochemistry, Ghent University, 9000 Ghent, Belgium; ⁸Bioinformatics Institute Ghent, Ghent University, 9052 Zwijnaarde, Belgium

Corresponding author: wout.bittremieux@uantwerpen.be, kris.laukens@uantwerpen.be

Abstract

As mass spectrometry-based proteomics has matured during the past decade a growing emphasis has been placed on quality control. For this purpose multiple computational quality control tools have been introduced. These tools generate a set of metrics that can be used to assess the quality of a mass spectrometry experiment. Here we review which different types of quality control metrics can be generated, and how they can be used to monitor both intra- and inter-experiment performance. We discuss the principal computational tools for quality control and list their main characteristics and applicability. As most of these tools have specific use cases it is not straightforward to compare their performance. For this survey we used different sets of quality control metrics derived from information at various stages in a mass spectrometry process and evaluated their effectiveness at capturing qualitative information about an experiment using a supervised learning approach. Furthermore, we discuss currently available algorithmic solutions that enable the usage of these quality control metrics for decision-making.

1 Introduction

In the past decade mass spectrometry (MS)-based proteomics has evolved into an extremely powerful analytical technique to identify and quantify proteins in complex biological samples. This high-throughput approach can yield a considerable volume of complex data for each experiment. As it has matured, over the last few years a growing emphasis has been placed on quality assurance (QA). This attention on quality assurance is of the utmost importance to safeguard confidence in the acquired results: in cases where this has been lacking mass spectrometry proteomics has sometimes suffered from exaggerated claims [4, 20]. To anticipate this evolution, a shift to “quality by design” is now taking place [51]. This means that the “designing and developing formulations and manufacturing processes ensure a predefined product quality.” As such, quality assurance consists of multiple aspects of which quality control (QC) is an essential component, but other elements such as a careful experimental design [12, 25, 35] are equally vital.

Whereas the experimental design has to be established prior to the initiation of an experiment, quality control takes place while or after the experimental results are obtained. Nonetheless, quality control and experimental design should not be discussed in isolation, as they are interwoven. For example, a QC sample can consist of a single peptide, a single protein digest, or a complex lysate, and this decision influences the type of QC metric(s) that can be investigated [5, 30, 41]. Furthermore, one has to decide how many QC runs to include in the experiment and to what extent and in which order these QC runs are interleaved with the biological samples under consideration. The goal

of quality control is then to leverage the experimental set-up to comprehend how well an instrument performs and how confident the results from the experiments are.

Related to the experimental design and based on the type of performance we want to monitor there are multiple approaches to quality control. A typical example consists of the use of QC samples with a simple sample content interleaved between the biological samples. The interesting aspect of such QC samples is that they have a controlled, limited, and known sample content. They are typically measured on a frequent basis, which allows to extract periodic information on the performance of the mass spectrometer. Of course, to understand this performance expressive QC metrics that provide information indicative of the quality of the experimental results need to be derived. Some straightforward and commonly used QC metrics include the number of identifications or the sequence coverage. Although these metrics give a global view of the performance, they do not allow us to pinpoint specific elements of the workflow where a failure might have arisen. Instead, more granular QC metrics providing information on the chromatography, the ion signal, the spectrum acquisition, etc., might be used.

Over the years dozens of QC metrics have been proposed, generated by a range of bioinformatics tools. In this paper we will list the main QC tools and explain their use cases and capabilities. Furthermore, we will provide an empirical assessment of which type of QC metrics is most adequate in detecting low-quality experiments.

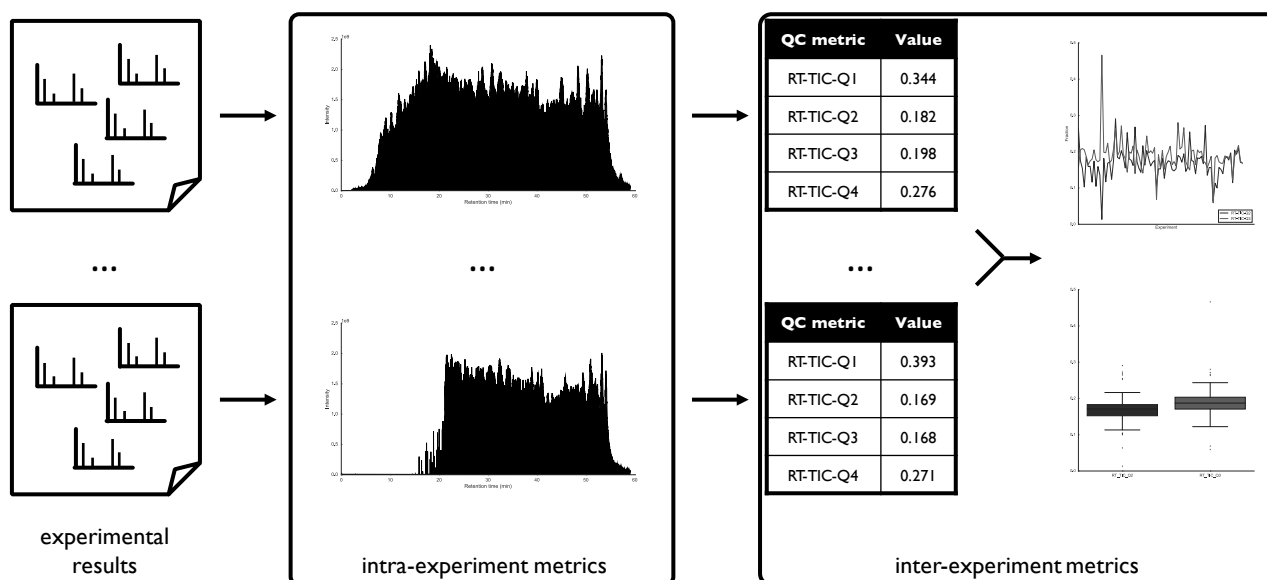


Figure 1: Intra-experiment metrics evaluate the quality of a single experiment, whereas inter-experiment metrics can be used to compare the quality of multiple experiments.

1.1 Quality control metrics

We can primarily distinguish QC metrics based on whether they represent information about a single experiment, or about multiple experiments, as illustrated in figure 1.

Intra-experiment metrics give information about a single experiment and are computed at the level of individual scans or identifications. These metrics show the evolution of a specific measure over the experiment run time, such as, for example a chromatogram of the total ion current (TIC) over the retention time (RT), or the mass accuracy of the identified spectra.

Inter-experiment metrics, on the other hand, assess a specific part of the quality of an experiment using a single measurement for the whole experiment. These values can subsequently be compared for multiple experiments, for example through a longitudinal analysis to evaluate the performance over time. Often an intra-experiment metric can be converted to an inter-experiment metric through summarization. This is illustrated in figure 1, where a TIC chromatogram enables the assessment of the chromatographic performance by visualizing the intensity distribution over the retention time. Using summary statistics this continuous information can be converted to inter-experiment metrics detailing the fraction of the total retention time that was required to accumulate a certain amount of the TIC, which gives a high-level assessment at the experiment level of the chromatographic stability.

To compare inter-experiment metrics multiple observations for different experiments are required. Therefore, QC tools that analyze these metrics usually include a database back-end for the persistent storage of historical data. On the other hand, intra-experiment metrics can be computed from only a single experiment and there is no comparison with external data. As a result, QC tools that exclusively generate intra-experiment metrics are generally easier to set up, as no external data storage needs to be provided. Because the use cases and requirements differ between these two

types of tools, we will further make a distinction between tools that generate metrics for individual experiments, tools that compare a limited group of experiments and do not necessarily require a complex back-end for data storage, and tools for longitudinal tracking that store QC data for a large number of experiments.

A second distinction between various metrics can be made based on from which stage in a mass spectrometry workflow they represent the quality of the system. As shown in figure 2, we can distinguish between instrument metrics, identification (ID)-free metrics, and ID-based metrics.

ID-free metrics and ID-based metrics are similar in the sense that they are both computed from the spectral results. *ID-free* metrics are derived solely from the spectral results, i.e. from the raw spectral data directly generated by the mass spectrometer. These metrics aim to capture information over the whole mass spectrometry workflow and include for example the shape of the peaks or the course of TIC detailing the chromatography, the number of MS1 and MS2 scans or the scan rate detailing the spectrum acquisition, or the charge state distribution detailing the ionization. The advantage of ID-free metrics is that they are generated directly from the raw spectral data, which makes it possible to instantly generate these metrics as soon as a mass spectrometry run has been completed.

ID-based metrics are derived from the spectral results as well, but they combine these data with subsequently obtained identification results. Examples include aforementioned metrics such as the number of identifications in terms of peptide-spectrum matches (PSMs), peptides, or proteins; or the sequence coverage for a known sample. Other detailed metrics can be computed as well, for example by comparing the difference in RT for similar identifications to assess the chromatographic stability, the number of spectra identified as the same peptide to measure the dynamic sampling, or by linking information similar to the ID-free

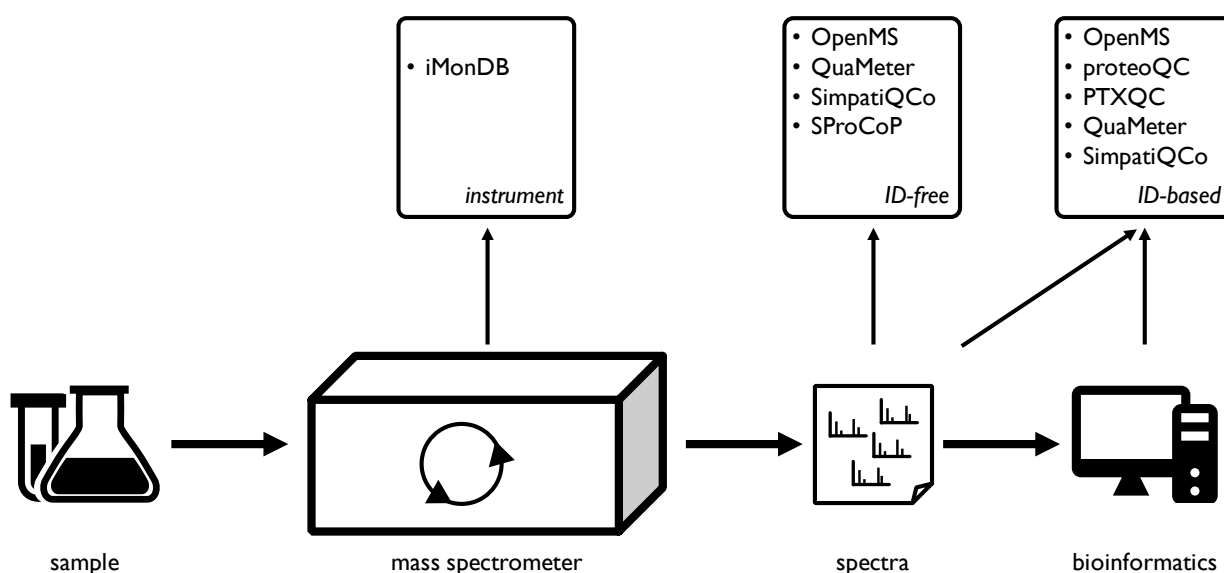


Figure 2: QC tools can capture qualitative information at different stages of a mass spectrometry experiment. For each type of QC metrics the representative tools are listed.

metrics with the identification results. Compared to ID-free metrics, the computation of ID-based metrics is somewhat more involved because it additionally requires the identifications results. Furthermore, the computation of ID-based metrics can be negatively influenced by suboptimal identification settings. However, in general the inclusion of identifications can provide a more detailed qualitative assessment of the experimental results.

Finally, *instrument* metrics do not look at the spectral data but derive information directly from instrument readouts. These are typically very sensitive, low-level metrics, such as the status of the ion source, the vacuum, or a turbo pump, depending on the type of instrument. An advantage of instrument metrics is that they directly indicate which part of the instrument is outside its normal range of operation. This facilitates troubleshooting and can be a driver for maintenance scheduling. On the other hand, these metrics cannot be directly related to the experimental results, instead they provide a secondary source of QC information. Furthermore, instrument metrics are instrument- and vendor-specific, and are typically not included in open file formats such as the mzML format [38].

Each distinct type of metric can give a different view on the quality of the data. However, not all metrics are always applicable; often metrics are especially relevant for a particular type of sample. For example, monitoring the sequence coverage is mostly applicable when using samples that contain a single protein digest, whereas the number of protein identifications is applicable to samples that consist of a complex lysate. Additionally, the type of experiment also plays an important role. For example, the number of identifications is very relevant for a discovery experiment, but less so for a targeted experiment. In contrast, instrument metrics are largely agnostic to the type of experiment and the sample content, but they can significantly vary between different instrument models and vendors.

2 QC tools

In recent years, quality control has become a key focus of attention in academic, industrial, and governmental proteomics laboratories. This trend is exemplified (and possibly driven) by the numerous QC tools that have been developed over the past few years. Initial work by Rudnick et al. [45] described for the first time how computational QC metrics can be used to objectively assess the quality of a mass spectrometry proteomics experiment. Whereas previously quality control was mostly performed manually by monitoring a few key measurements, this work showed how a comprehensive set of QC metrics can be used to thoroughly investigate the system performance. A set of 46 mainly ID-based metrics was defined and implemented in a pipeline of Perl programs by researchers at the National Institute of Standards and Technology (NIST), called NIST MSQC. This set of metrics has since then been reimplemented in several lab-specific data processing pipelines. Support for NIST MSQC itself has been discontinued in early 2016 and the original implementation is no longer available, but several of the reimplementations remain under active development.

It has been demonstrated that computational QC metrics provide objective criteria that can accurately capture the quality of a mass spectrometry experiment, and there has been a proliferation of tools that can compute such metrics. Here, we will detail the primary tools, their characteristics, and their usage. Table 1 provides an overview of the discussed tools.

2.1 Tools evaluating individual experiments

2.1.1 QuaMeter

QuaMeter was initially developed as a user-friendly and open-source alternative to NIST MSQC. NIST MSQC con-

Tool	Interface	Operating system	Experiment type	Instrument	ID-free	ID-based	Website
QuaMeter [33, 59]	command-line	Windows, Linux	discovery DDA	×	✓	✓	http://proteowizard.sourceforge.net/
OpenMS [58]	KNIME	cross-platform	discovery DDA	×	✓	✓	http://www.openms.de/
proteoQC [60]	R	cross-platform	discovery DDA	×	×	✓	http://bioconductor.org/packages/proteoQC
PTXQC [8]	R	Windows, cross-platform	discovery & quantification DDA	×	×	✓	https://github.com/cbielow/PTXQC
SProCoP [6]	Skyline	Windows	discovery, targeted SRM & PRM	×	✓	×	http://proteome.gs.washington.edu/software/skyline/tools/sprocop.html
SimpatiQCo [43]	web	Windows	discovery DDA	×	✓	✓	http://ms.imp.ac.at/?goto=simpatiqco
iMonDB [10]	GUI	Windows	any	✓	×	×	https://bitbucket.org/proteinspector/imondb/

Table 1: An overview of the discussed QC tools and their main characteristics.

sisted of a graphical user interface (GUI) wrapper around multiple individual tools and scripts with various inter-dependencies, which resulted in a complex pipeline. Additionally, some elements of this pipeline could only be modified to a limited extent. NIST MSQC could exclusively compute metrics from Thermo Scientific raw files, and only supported three search engines to provide identifications: the NIST MSPepSearch or the SpectraST [31] spectral library search engines, or the OMSSA [24] sequence database search engine. These limitations restricted the applicability of NIST MSQC. Instead, QuaMeter consists of a single multi-platform command-line application that is able to compute QC metrics from raw files originating from instruments produced by multiple vendors. Using the ProteoWizard [13] library it is able to read spectral data stored in a wide variety of vendor-specific raw files (restricted to the Windows platform) and open standard file formats, such as mzML [38]. Further, it can utilize identification results produced by any search engine in the standard mzIdentML [26] or pepXML format through external processing using IDPicker [32].

The initial QuaMeter version [33] computed a set of 42 ID-based QC metrics equivalent to those defined by Rudnick et al. [45]. In a subsequent version QuaMeter improved upon this by also including functionality to compute a set of 45 ID-free QC metrics [59]. Both sets of metrics are inter-experiment summary metrics, although the output is exported to simple tab-delimited text files, so the visualization and analysis thereof has to be done using external software or code scripts. Without advanced visualization or analysis functionality QuaMeter focuses solely on computing QC metrics. Especially the set of ID-free metrics, which requires only the spectral data, can very easily be computed. For the set of ID-based metrics some prior processing of the identification results by IDPicker is required, which can make this process slightly more cumbersome. Only a limited configuration is required, and through the

command-line functionality the computation can easily be automated. This makes QuaMeter a powerful tool that computes an extensive set of inter-experiment QC metrics.

2.1.2 OpenMS

OpenMS is a comprehensive open-source software library that offers a wide range of algorithms and tools for mass spectrometry-based proteomics and metabolomics [49]. It consists of various small processing tools that can be used to construct complex analysis workflows [2, 27]. These workflows can be designed visually using the KNIME workflow engine [7], where each tool functions as an individual node in the workflow.

The various OpenMS nodes can be used to build complex QC pipelines [58]. The provided QC nodes can compute a set of intra-experiment metrics, consisting of both ID-free and ID-based metrics. OpenMS supports a range of search engines to generate identifications for the ID-based metrics, for which there exist specific nodes, including Mascot, MS-GF+ [28], Myrimatch [52], OMSSA [24], and X!Tandem [16]. Example QC metrics include the number of spectra (identified or otherwise), peptides, and proteins; mass accuracy statistics; and the mass over charge and retention time acquisition ranges. These metrics are complemented by various plots that provide further details, such as a TIC chromatogram, a histogram of the mass accuracy of the identified peptides, or a histogram of the charge distribution of the detected ion features. OpenMS exports this information to an Extensible Markup Language (XML) qcML file [58], which can be visualized in a web browser through an embedded stylesheet, or to a portable document format (PDF) report.

Due to the wealth of algorithms and tools that are available in the OpenMS software library, the provided QC workflows can potentially be easily extended to compute additional metrics. Furthermore, there is no need to be restricted

to algorithms natively provided by OpenMS, as the available functionality can easily be extended through custom nodes, for example by using the built-in support for the R statistical programming language [44]. This makes it possible to build granular workflows and achieve a very fine-grained control, although expert knowledge of the OpenMS ecosystem and the KNIME environment is recommended to do so. The constructed workflows can subsequently be exported and shared. Both OpenMS and KNIME are cross-platform tools, ensuring the universal applicability of these workflows.

2.2 Tools comparing groups of experiments

2.2.1 proteoQC

The proteoQC package [60] for the R programming language [23, 44] can be used to generate a Hyper Text Markup Language (HTML) report detailing the experimental quality. Prior to executing proteoQC the experimental design has to be specified by configuring each spectral data file representing a sample as belonging to a specific fraction, technical replicate, and biological replicate. The generated QC report contains intra-experiment metrics for each individual sample, as well as aggregated information to compare samples at the level of their fractions, technical replicates, and biological replicates.

To generate a set of intra-experiment ID-based metrics for each sample, proteoQC uses the rTANDEM package [22] to interface the X!Tandem [16] sequence database search engine in R to provide identification results. For each sample some individual metrics and QC plots are generated, such as a breakdown of the precursor ion charge states, the mass accuracy, information on the number of spectra and peptides that were used to identify distinct proteins during protein inference, etc. Furthermore, when identifying the data proteoQC automatically adds the common Repository of Adventitious Proteins (cRAP) (<http://www.thegpm.org/crap/>) database to the user-provided protein database. The cRAP database contains contaminants such as common laboratory proteins, like trypsin, or contaminants transferred through dust or contact, like keratin, and proteoQC reports which of these contaminants were detected in the samples. Additionally, proteoQC reports on the reproducibility of the results by comparing the number of identified spectra, peptides, and proteins per fraction, technical replicate, and biological replicate, and their overlap between the replicates.

By incorporating the experimental design proteoQC can make informed comparisons between individual samples, which provides QC information on an additional level. Furthermore, proteoQC is fully cross-platform within the popular R programming language. However, as the QC pipeline has to be configured programmatically, some R experience is recommended to utilize proteoQC.

2.2.2 PTXQC

Proteomics Quality Control (PTXQC) [8] is an R-based quality control pipeline for MaxQuant [15], a highly popular software suite for quantitative proteomics. Like MaxQuant, PTXQC supports a wide range of quantitative proteomics workflows, including stable isotope labeling with amino acids in cell culture (SILAC), tandem mass tags (TMT), and

label-free quantification. After initial processing of the spectral data by MaxQuant, PTXQC uses the MaxQuant output results to compute various QC metrics. PTXQC requires as input the custom text files generated by MaxQuant and the MaxQuant configuration settings, and hence cannot be used to process any other type of data. As PTXQC is written in the R programming language, it is fully cross-platform. Additionally, easy drag-and-drop functionality to execute the QC analyses is provided for the Windows operating system.

PTXQC produces an extensive report that contains a set of 24 intra- and inter-experiment metrics. These metrics are divided into four categories corresponding to the specific MaxQuant output source the metrics are derived from: “ProteinGroups”, “Evidence”, “Msms”, and “MsmsScans”. The metrics cover a wide range of information, including the intensity of the detected features and peptides, the potential presence of contaminants, the mass accuracy of the identified peptides and fragments, the number of missed cleavages detailing the enzyme specificity, and the number of identified peptides and proteins. Other metrics are specifically related to the MaxQuant “match-between-runs (MBR)” [14] functionality. MBR aligns the retention times of multiple runs and transfers their identifications across features that have the same accurate mass and a similar retention time, providing more data for the downstream quantification of proteins. PTXQC assesses the MBR performance by evaluating the retention time alignment and by checking whether the identification transfer seems correct. All of these metrics are then visualized and compared between the different raw files that constitute the considered MaxQuant project using detailed figures. Furthermore, each of the metrics is converted to an individual score for each experiment using automated scoring functions. Most of these scores are absolute scores generated by comparing the observation to a threshold, for example such as whether the number of detected contaminants is too excessive, or generated by evaluating a specific characteristic of the observation, for example such as the extent to which the mass deviations are centered around zero. Other scores are computed for a single raw file using the other raw files as a reference, for example by comparing the number of missed cleavages in each individual raw file to the average number of missed cleavages. Finally, some other scores are evaluated relative to settings extracted from MaxQuant, such as the mass accuracy compared to the width of the precursor mass window. All these scoring functions generate inter-experiment metrics that are used to compare the quality of the different experiments. Usefully, PTXQC provides a heatmap overview of the inter-experiment metrics, which yields an assessment of the quality at a glance and facilitates pinpointing the low-performing experiments.

Although PTXQC can exclusively be used to analyze MaxQuant results, through this tight integration it is able to compute some highly relevant and specialized QC metrics. These metrics do not only assess the quality of the spectral data, but also provide information on the subsequent bioinformatics processing by MaxQuant. Furthermore, the addition of a high-level heatmap at the start of the report is very useful to get a quick overview of the quality, after which the more detailed visualizations can be employed to

further investigate potential problems.

2.2.3 SProCoP

Statistical Process Control in Proteomics (SProCoP) [6] is a QC script written in R [44] that can be used as a plugin [11] for the popular Skyline [34] tool for targeted proteomics. SProCoP applies well-established statistical process control techniques such as the Shewhart control chart and the Pareto chart. The purpose of a Shewhart control chart is to track performance over time and identify outliers that deviate excessively from the expected behavior. Further, the Pareto chart is a combination of a bar and line graph, which displays the number of deviating measurements for each metric along with its cumulative percentage, and provides feedback on which metrics are more variable and may require attention.

Using these statistical process control techniques SProCoP monitors the performance of five inter-experiment QC metrics based on targeted peptides present in QC samples with a known sample content or spiked into real samples: signal intensity, mass measurement accuracy, retention time reproducibility, peak full width at half maximum (FWHM), and peak symmetry. Measurement thresholds are defined empirically based on a reference set of samples with a known good quality, after which the performance of other samples in the Skyline project can be investigated.

Through its integration with Skyline SProCoP is vendor-independent and can be used for a wide range of targeted and discovery workflows. Additionally these statistical process control techniques are available online (<http://www.qcmylcms.com/>) and have been implemented in the Panorama [47] repository for targeted proteomics from Skyline. Panorama AutoQC is a utility application that monitors for new data files and automatically invokes Skyline to process the data. The QC metrics are stored in Panorama and the statistical process control charts similar to SProCoP can be visualized through the Panorama web application.

2.3 Tools for longitudinal tracking

2.3.1 SimpatiQCo

SIMple AuTomatIc Quality COnTrol (SimpatiQCo) [43] not only computes various QC metrics, it also stores and visualizes these metrics for longitudinal monitoring of the system performance. It uses a PostgreSQL database as back-end, and an Apache webserver to provide a web-based front-end for configuration and visualization.

SimpatiQCo can compute QC metrics from a limited selection of Thermo Scientific and SCIEX instruments. Raw files from these instruments can be uploaded to the web server manually, or can be added automatically through a “hot folder” that is monitored continuously for new raw files. These raw files are then submitted to a linked Mascot server for peptide identifications. Next, SimpatiQCo calculates a range of ID-free and ID-based QC metrics such as the number of MS1 and MS2 scans, the number of identified PSMs and proteins, the TIC, and information on lock masses (if applicable). Further, specific peptides and proteins can be investigated in detail using metrics such as the peak area and width and the elution time of peptides of interest, and

the protein sequence coverage. For each QC metric the range of acceptable values is learned based on the historical observations using robust statistical measures to take outlying values into account. This information is then displayed in the metric plots using a color-coded background band to highlight deviating system performance. Further, external messages can be entered manually, for example pertaining to instrument maintenance. These messages will be superimposed on the metric plots to relate the external events to the evolution of the metrics.

SimpatiQCo consists of a number of different components, such as the database, the web server, and various processing tools. These components need to be installed individually, and although a step-by-step installation guide is available online, this complicated process is not recommended for novice users. Furthermore, not all of the configuration can be done through the graphical web-based client. For example, to process raw files these must be able to be linked to a specific instrument. Unfortunately, an instrument definition can only be created by manually adding a record in the corresponding table of the PostgreSQL database.

SimpatiQCo is a powerful tool to track system performance over time, albeit with some technical limitations. Namely, SimpatiQCo is only able to process raw files generated on a limited number of instrument models and only supports the commercial Mascot search engine for peptide identifications.

2.3.2 iMonDB

Unlike the previous tools the Instrument MONitoring DataBase (iMonDB) [10] does not compute metrics from the spectral results, but extracts instrument metrics from the raw files. The iMonDB uses a MySQL database to store its information. This database acts as a server, with two separate standalone GUI applications that can connect to the database as clients, each with a specific task: the iMonDB Collector processes raw files and stores the instrument metrics in the database, whereas the iMonDB Viewer retrieves the information from the database and visualizes it.

The iMonDB supports a wide range of instruments manufactured by Thermo Scientific, although it does not support other instrument vendors. Prior to extracting instrument metrics from a raw file, a corresponding instrument definition has to be created. This can be done through the iMonDB Collector, which allows the full configuration through its GUI. Further, extraction of the instrument metrics can be done manually through the GUI, or can be done through command-line functionality provided by the iMonDB Collector. This command-line functionality can be used to automatically run the iMonDB Collector using an external scheduling tool, such as the native operating system scheduler.

The behavior over time of the metrics for each instrument can be viewed using the iMonDB Viewer. Similar to functionality provided by SimpatiQCo it is possible to add additional information pertaining to external events and show this on the metric plots to link this to the evolution of the metrics. It is also possible to export a PDF file of the external events for reporting purposes.

A unique aspect of the iMonDB is that this is the only

tool that is able to systematically analyze instrument metrics. The advantage of these instrument metrics, which provide information at the lowest level, is their high sensitivity, which makes it possible to detect emerging defects in a timely fashion. However, because these metrics are instrument-dependent they are usually not retained during conversion to open formats, such as mzML [38]. Due to this limitation the iMonDB needs to work with vendor-specific raw files directly, which is currently limited to Thermo Scientific raw files. Furthermore, there is a multitude of instrument metrics that are extracted, which makes it hard to comprehend which metrics are most useful to monitor systematically, even for expert users. Nevertheless, these instrument metrics can be very useful to detect malfunctioning instrument elements before these have a deleterious effect on the experimental results, preventing potential loss of valuable sample content.

2.4 Other tools

As mentioned previously, NIST MSQC [45] was the first tool that generated computational QC metrics, although it was recently retired in early 2016.

Metriculator [54] is a web-based tool for storing and visualizing QC metrics longitudinally. However, Metriculator does not compute QC metrics directly but critically depends upon an embedded version of NIST MSQC. Unfortunately, the installation process for Metriculator is not very straightforward; it has many Ruby dependencies whose installation might fail, and which are presently outdated or even no longer supported.

LogViewer [50] is a simple visualization tool that presents a set of 11 instrument metrics, such as MS1 and MS2 ions injection times, and ID-free metrics, such as the charge state and mass distributions. As input it uses log files from Thermo instruments exported by RawXtract [39], which has been deprecated presently.

A different approach is used by SprayQc [46]. Whereas the other discussed tools compute QC metrics post-acquisition, SprayQc directly interfaces with peripheral equipment to continuously monitor its performance. SprayQc is able to automatically track the stability of the electrospray through computer vision, the status of the liquid chromatography (LC) pumps, the temperature of the column oven, and the continuity of the data acquisition. In case a malfunctioning is detected SprayQc can automatically take corrective actions and warn the instrument operator. This is a valuable approach to minimize the loss of precious sample content and provide early notifications, and it can complement the other QC tools that provide a post-acquisition quality assessment.

3 Metrics evaluation

We compared various sets of metrics to assess their effectiveness in expressing the quality of a mass spectrometry proteomics experiment. Typically this is not a straightforward task because, as we have reviewed in the previous sections, each QC tool has its own characteristics and requirements, and use cases can vary as some tools are spe-

cific to certain experimental workflows and sample types. Meanwhile most tools also represent some of their QC information through visualizations. Although these quickly provide useful insights for human users, this data is not suitable for an objective, automatic comparison.

To compare different types of metrics we used the set of instrument metrics computed by the iMonDB [10], the set of ID-free metrics computed by QuaMeter [59], and the set of ID-based metrics as identified by Rudnick et al. [45]. These sets of metrics are very comprehensive and all of these inter-experiment metrics can readily be used to compare experiments to each other. To be able to determine whether or not these metrics can capture qualitative information about an experiment, we used a public dataset for which the quality of the experiments is known. The dataset consists of a number of complex quality control LC-MS runs performed on several different instruments at the Pacific Northwest National Laboratory (PNNL) [3]. Each sample had an identical content (whole cell lysate of *Shewanella oneidensis*), and the quality of the various runs has been manually annotated by expert instrument operators as being either “good”, “ok”, or “poor”. We split up the various runs depending on the instrument type, being either “Exactive”, “LTQ IonTrap”, “LTQ Orbitrap”, or “Velos Orbitrap”, with each of these instrument groups consisting of multiple individual instruments. We refer to the original publication by Amidan et al. [3] for further information on the experimental procedures and the dataset details.

This public dataset already contains the precomputed set of ID-free metrics by QuaMeter and the set of ID-based metrics by SMAQC (the PNNL in-house reimplementation of the NIST MSQC metrics defined by Rudnick et al. [45]; <https://github.com/PNNL-Comp-Mass-Spec/SMAQC>). We further used the iMonDB to compute the set of instrument metrics. To this end all experimental raw files, pre-computed QC metrics, and the expert quality annotations were retrieved from the PRoteomics IDentifications (PRIDE) database [57].

To quantify the expressiveness of these three sets of metrics, each capturing a different type of QC information, we employed a binary classifier. As the quality of the experiments was manually assessed by expert instrument operators, this labeling can be used as the ground truth to train the classifier. We used the acceptable experiments, with their quality designated as either “good” or “ok”, as the positive class, and the inferior experiments, with their quality designated as “poor”, as the negative class. When given an experiment represented by its QC metrics, the classification task consists of correctly predicting the experiment’s quality. Prior to training the classifier we removed redundant features that have a very low variance and we rescaled the features robust to outliers by centering by the median and scaling by the interquartile range. Next, for each separate instrument type we trained a random forest classifier, for which we split the data into 65%–35% training and testing subsets that are equally stratified according to their quality labels. This classifier has been coded in Python and uses the random forest implementation from scikit-learn [42], along with functionality provided by NumPy [56] and pandas [40]. The code is available as open source at <https://bitbucket.org/proteinspector/qc-evaluation/>.




Figure 3: ROC curve showing the classification performance for the Velos Orbitrap instrument type. ROC curves for the other instrument types indicate similar results (data not shown).

As illustrated by the receiver operator characteristic (ROC) curve in figure 3 all three types of QC metrics are adept at discriminating high-quality experiments from low-quality experiments. This shows that all of the different tools can give us valuable insights into the quality of an experiment, and that information captured at various different stages of the mass spectrometry process should be investigated. ID-based metrics slightly outperform ID-free metrics, most likely because the ID-based metrics can employ additional information provided by the identifications. This difference is minimal however, which is perhaps not surprising as both types of metrics take similar properties of the spectra into account. This reinforces previous research which showed that ID-based metrics are not significantly influenced by slight differences in the identifications, such as when using an alternative search engine [33]. This also shows the excellent efficacy of ID-free metrics in objectively evaluating the quality based solely on spectral information. Because ID-based metrics require additional computational steps to obtain the identifications, whereas ID-free metrics can be directly computed from the spectral results, ID-free metrics might be preferred if a speedy quality assessment is required. In contrast, instrument metrics perform a little worse at correctly identifying low-quality experiments. This is likely because they are only secondary results that are not always directly related to the data quality. Nevertheless, these metrics still have merit as they do not depend on a specific type of experiment or sample content, but are applicable on all occasions. Furthermore, by combining the individual classifiers for the various types of metrics in an ensemble classifier a further performance gain can be achieved because the different types of metrics each provide a complementary view on the quality.

4 Using QC metrics for decision-making

As tools for computational quality control have proliferated in recent years, the challenge in this field is now shifting from the computation of QC metrics toward informed

decision-making based on these metrics. However, interpreting these metrics is not trivial. First, considerable domain knowledge is required to understand what each metric signifies. Second, the metrics form a high-dimensional data space, which complicates their analysis. Different elements in a mass spectrometry workflow do not function in isolation but instead influence each other, which has to be taken into account while analyzing metrics representing information about these elements. Therefore, univariate approaches are generally insufficient; instead multivariate approaches that can deal with the high-dimensional data space should be preferred, while also taking the curse of dimensionality into account [1].

To this end Wang et al. [59] have developed a robust multivariate statistical toolkit to interpret QC metrics. They have used a principal component analysis (PCA) transformation to reduce the data to a low-dimensional approximation, in which they were able to successfully detect outlier low-quality experiments based on pairwise dissimilarities. Furthermore, they developed an analysis of variance (ANOVA) model which enabled them to identify whether the observed variability was attributable to lab-dependent factors, batch effects, or biological variability. Such work driving the understanding of QC metrics is highly valuable, and these analyses have been applied to great effect for multiple studies. For example, it was used to assess the quality of the experimental results for various studies conducted by the National Cancer Institute (NCI) Clinical Proteomic Tumor Analysis Consortium (CPTAC) [48, 53, 61].

Similar work was done by Bittremieux et al. [9], who applied unsupervised outlier detection to identify low-quality experiments. Subsequently they used a specialized outlier interpretation technique to determine which QC metrics mostly contributed to the decrease in quality. The advantage of this approach is that all QC metrics are used to identify low-quality experiments, unlike when using a dimensionality reduction, such as PCA, which discards some of the information. Meanwhile, the advanced outlier interpretation pinpointing the most relevant QC metrics can yield actionable information for domain experts to optimize their experimental set-up.

Whereas these previous analyses used unsupervised techniques, Amidan et al. [3] trained a supervised classifier to discriminate low-quality experiments from high-quality experiments. A supervised approach will generally perform better than an unsupervised approach but will require initial training. Furthermore, a supervised classifier might have to be retrained to adapt it to data generated by a different instrument or in a different laboratory. Amidan et al. [3] have expended significant effort in manually annotating the quality of over a thousand experiments to generate training data, which allowed them to build a highly performant logistic regression classifier.

These analyses are extremely valuable, as they allow us to achieve a deeper understanding of the mass spectrometry processes and the properties of what makes a high-quality experiment. These algorithmic approaches provide a thorough quality assessment of the spectral data, which enforces informed decision-making, and which has the potential to automatically drive the spectral acquisition in the future.

5 Conclusion

We have given an overview of the available computational tools to generate QC metrics for mass-spectrometry-based proteomics. These tools enable assessing the performance of the experimental set-up and detecting unreliable results. These are essential requirements to inspire confidence in the experimental results, which will prove to be a crucial step in the maturation of proteomics technologies, and which will allow us to for example routinely apply these technologies into a clinical setting [36, 51]. Another potential application where an accurate assessment of the data quality is paramount, is in the reuse of public data [19, 21, 29, 37]. As public data repositories keep expanding and the potential for data reuse grows, we envision that data submissions to public repositories will soon have to be accompanied by QC parameters at the time of submission, or will have a standard set of QC metrics calculated automatically after submission [37].

Finally, most current QC tools are limited to the typical use case of bottom-up data-dependent acquisition (DDA) discovery experiments, and their QC metrics often cannot be directly translated to other types of experiments. Less research has been done on QC for other types of workflows, such as data-independent acquisition (DIA) [18] or top-down proteomics [55], or even related mass spectrometry-based domains, such as metabolomics [17]. In the next few years we will likely see the efforts on QC expanded to these types of workflows as well, which will further bolster the diverse and powerful mass spectrometry ecosystem.

Acknowledgments

W.B., D.V., L.M., and K.L. acknowledge the support of the VLAIO grant “InSPECTor” (IWT project 120025).

References

- [1] Aggarwal, C. C., Hinneburg, A., and Keim, D. A. “On the Surprising Behavior of Distance Metrics in High Dimensional Space.” In: *Proceedings of the 8th International Conference on Database Theory - ICDT '01*. Vol. 1973. Lecture Notes in Computer Science. London, England: Springer Berlin Heidelberg, 2001, pp. 420–434. DOI: [10.1007/3-540-44503-X_27](https://doi.org/10.1007/3-540-44503-X_27).
- [2] Aiche, S., Sachsenberg, T., Kenar, E., Walzer, M., et al. “Workflows for Automated Downstream Data Analysis and Visualization in Large-Scale Computational Mass Spectrometry.” In: *PROTEOMICS* 15.8 (Apr. 2015), pp. 1443–1447. DOI: [10.1002/pmic.201400391](https://doi.org/10.1002/pmic.201400391).
- [3] Amidan, B. G., Orton, D. J., LaMarche, B. L., Monroe, M. E., et al. “Signatures for Mass Spectrometry Data Quality.” In: *Journal of Proteome Research* 13.4 (Apr. 4, 2014), pp. 2215–2222. DOI: [10.1021/pr401143e](https://doi.org/10.1021/pr401143e).
- [4] Baggerly, K. A., Morris, J. S., Edmonson, S. R., and Coombes, K. R. “Signal in Noise: Evaluating Reported Reproducibility of Serum Proteomic Tests for Ovarian Cancer.” In: *JNCI Journal of the National Cancer Institute* 97.4 (Feb. 16, 2005), pp. 307–309. DOI: [10.1093/jnci/dji008](https://doi.org/10.1093/jnci/dji008).
- [5] Bereman, M. S. “Tools for Monitoring System Suitability in LC MS/MS Centric Proteomic Experiments.” In: *PROTEOMICS* 15 (5–6 Mar. 2015), pp. 891–902. DOI: [10.1002/pmic.201400373](https://doi.org/10.1002/pmic.201400373).
- [6] Bereman, M. S., Johnson, R., Bollinger, J., Boss, Y., et al. “Implementation of Statistical Process Control for Proteomic Experiments via LC MS/MS.” In: *Journal of The American Society for Mass Spectrometry* 25.4 (Apr. 2014), pp. 581–587. DOI: [10.1007/s13361-013-0824-5](https://doi.org/10.1007/s13361-013-0824-5).
- [7] Berthold, M. R., Cebron, N., Dill, F., Gabriel, T. R., et al. “KNIME - the Konstanz Information Miner: Version 2.0 and Beyond.” In: *ACM SIGKDD Explorations Newsletter* 11.1 (June 2009), pp. 26–31. DOI: [10.1145/1656274.1656280](https://doi.org/10.1145/1656274.1656280).
- [8] Bielow, C., Mastrobuoni, G., and Kempa, S. “Proteomics Quality Control: Quality Control Software for MaxQuant Results.” In: *Journal of Proteome Research* 15.3 (Mar. 4, 2016), pp. 777–787. DOI: [10.1021/acs.jproteome.5b00780](https://doi.org/10.1021/acs.jproteome.5b00780).
- [9] Bittremieux, W., Meysman, P., Martens, L., Valkenburg, D., et al. “Unsupervised Quality Assessment of Mass Spectrometry Proteomics Experiments by Multivariate Quality Control Metrics.” In: *Journal of Proteome Research* 15.4 (Apr. 1, 2016), pp. 1300–1307. DOI: [10.1021/acs.jproteome.6b00028](https://doi.org/10.1021/acs.jproteome.6b00028).
- [10] Bittremieux, W., Willems, H., Kelchtermans, P., Martens, L., et al. “iMonDB: Mass Spectrometry Quality Control through Instrument Monitoring.” In: *Journal of Proteome Research* 14.5 (May 1, 2015), pp. 2360–2366. DOI: [10.1021/acs.jproteome.5b00127](https://doi.org/10.1021/acs.jproteome.5b00127).
- [11] Broudy, D., Killeen, T., Choi, M., Shulman, N., et al. “A Framework for Installable External Tools in Skyline.” In: *Bioinformatics* 30.17 (Sept. 1, 2014), pp. 2521–2523. DOI: [10.1093/bioinformatics/btu148](https://doi.org/10.1093/bioinformatics/btu148).
- [12] Cairns, D. A. “Statistical Issues in Quality Control of Proteomic Analyses: Good Experimental Design and Planning.” In: *PROTEOMICS* 11.6 (Mar. 2011), pp. 1037–1048. DOI: [10.1002/pmic.201000579](https://doi.org/10.1002/pmic.201000579).
- [13] Chambers, M. C., Maclean, B., Burke, R., Amodei, D., et al. “A Cross-Platform Toolkit for Mass Spectrometry and Proteomics.” In: *Nature Biotechnology* 30.10 (Oct. 10, 2012), pp. 918–920. DOI: [10.1038/nbt.2377](https://doi.org/10.1038/nbt.2377).
- [14] Cox, J., Hein, M. Y., Lubner, C. A., Paron, I., et al. “Accurate Proteome-Wide Label-Free Quantification by Delayed Normalization and Maximal Peptide Ratio Extraction, Termed MaxLFQ.” In: *Molecular & Cellular Proteomics* 13.9 (Sept. 1, 2014), pp. 2513–2526. DOI: [10.1074/mcp.M113.031591](https://doi.org/10.1074/mcp.M113.031591).

- [15] Cox, J. and Mann, M. "MaxQuant Enables High Peptide Identification Rates, Individualized p.p.b.-Range Mass Accuracies and Proteome-Wide Protein Quantification." In: *Nature Biotechnology* 26.12 (Nov. 30, 2008), pp. 1367–1372. DOI: 10.1038/nbt.1511.
- [16] Craig, R. and Beavis, R. C. "TANDEM: Matching Proteins with Tandem Mass Spectra." In: *Bioinformatics* 20.9 (Feb. 19, 2004), pp. 1466–1467. DOI: 10.1093/bioinformatics/bth092.
- [17] Dunn, W. B., Wilson, I. D., Nicholls, A. W., and Broadhurst, D. "The Importance of Experimental Design and QC Samples in Large-Scale and MS-Driven Untargeted Metabolomic Studies of Humans." In: *Bioanalysis* 4.18 (Sept. 2012), pp. 2249–2264. DOI: 10.4155/bio.12.204.
- [18] Egertson, J. D., MacLean, B., Johnson, R., Xuan, Y., et al. "Multiplexed Peptide Analysis Using Data-Independent Acquisition and Skyline." In: *Nature Protocols* 10.6 (May 21, 2015), pp. 887–903. DOI: 10.1038/nprot.2015.055.
- [19] Eisenacher, M., Schnabel, A., and Stephan, C. "Quality Meets Quantity - Quality Control, Data Standards and Repositories." In: *PROTEOMICS* 11.6 (Mar. 2011), pp. 1031–1036. DOI: 10.1002/pmic.201000441.
- [20] Ezkurdia, I., Calvo, E., Del Pozo, A., Vázquez, J., et al. "The Potential Clinical Impact of the Release of Two Drafts of the Human Proteome." In: *Expert Review of Proteomics* 12.6 (Nov. 2, 2015), pp. 579–593. DOI: 10.1586/14789450.2015.1103186.
- [21] Foster, J. M., Degroevé, S., Gatto, L., Visser, M., et al. "A Posteriori Quality Control for the Curation and Reuse of Public Proteomics Data." In: *PROTEOMICS* 11.11 (June 11, 2011), pp. 2182–2194. DOI: 10.1002/pmic.201000602.
- [22] Fournier, F., Joly Beuparlant, C., Paradis, R., and Droit, A. "rTANDEM, an R/Bioconductor Package for MS/MS Protein Identification." In: *Bioinformatics* 30.15 (Aug. 1, 2014), pp. 2233–2234. DOI: 10.1093/bioinformatics/btu178.
- [23] Gatto, L. and Christoforou, A. "Using R and Bioconductor for Proteomics Data Analysis." In: *Biochimica et Biophysica Acta (BBA) - Proteins and Proteomics* 1844.1 (Jan. 2014), pp. 42–51. DOI: 10.1016/j.bbapap.2013.04.032.
- [24] Geer, L. Y., Markey, S. P., Kowalak, J. A., Wagner, L., et al. "Open Mass Spectrometry Search Algorithm." In: *Journal of Proteome Research* 3.5 (Oct. 11, 2004), pp. 958–964. DOI: 10.1021/pr0499491.
- [25] Hu, J., Coombes, K. R., Morris, J. S., and Baggerly, K. A. "The Importance of Experimental Design in Proteomic Mass Spectrometry Experiments: Some Cautionary Tales." In: *Briefings in Functional Genomics and Proteomics* 3.4 (Jan. 1, 2005), pp. 322–331. DOI: 10.1093/bfpgp/3.4.322.
- [26] Jones, A. R., Eisenacher, M., Mayer, G., Kohlbacher, O., et al. "The mzIdentML Data Standard for Mass Spectrometry-Based Proteomics Results." In: *Molecular & Cellular Proteomics* 11.7 (July 1, 2012), pp. M111.014381–M111.014381. DOI: 10.1074/mcp.M111.014381.
- [27] Junker, J., Bielow, C., Bertsch, A., Sturm, M., et al. "TOPPAS: A Graphical Workflow Editor for the Analysis of High-Throughput Proteomics Data." In: *Journal of Proteome Research* 11.7 (July 6, 2012), pp. 3914–3920. DOI: 10.1021/pr300187f.
- [28] Kim, S. and Pevzner, P. A. "MS-GF+ Makes Progress towards a Universal Database Search Tool for Proteomics." In: *Nature Communications* 5 (Oct. 31, 2014), p. 5277. DOI: 10.1038/ncomms6277.
- [29] Kinsinger, C. R., Apffel, J., Baker, M., Bian, X., et al. "Recommendations for Mass Spectrometry Data Quality Metrics for Open Access Data (Corollary to the Amsterdam Principles)." In: *Molecular & Cellular Proteomics* 10.12 (Feb. 3, 2012), O111.015446–O111.015446. DOI: 10.1074/mcp.O111.015446.
- [30] Köcher, T., Pichler, P., Swart, R., and Mechtler, K. "Quality Control in LC-MS/MS." In: *PROTEOMICS* 11.6 (Mar. 6, 2011), pp. 1026–1030. DOI: 10.1002/pmic.201000578.
- [31] Lam, H., Deutsch, E. W., Eddes, J. S., Eng, J. K., et al. "Development and Validation of a Spectral Library Searching Method for Peptide Identification from MS/MS." In: *PROTEOMICS* 7.5 (Mar. 5, 2007), pp. 655–667. DOI: 10.1002/pmic.200600625.
- [32] Ma, Z.-Q., Dasari, S., Chambers, M. C., Litton, M. D., et al. "IDPicker 2.0: Improved Protein Assembly with High Discrimination Peptide Identification Filtering." In: *Journal of Proteome Research* 8.8 (Aug. 7, 2009), pp. 3872–3881. DOI: 10.1021/pr900360j.
- [33] Ma, Z.-Q., Polzin, K. O., Dasari, S., Chambers, M. C., et al. "QuaMeter: Multivendor Performance Metrics for LC-MS/MS Proteomics Instrumentation." In: *Analytical Chemistry* 84.14 (July 17, 2012), pp. 5845–5850. DOI: 10.1021/ac300629p.
- [34] MacLean, B., Tomazela, D. M., Shulman, N., Chambers, M., et al. "Skyline: An Open Source Document Editor for Creating and Analyzing Targeted Proteomics Experiments." In: *Bioinformatics* 26.7 (Apr. 1, 2010), pp. 966–968. DOI: 10.1093/bioinformatics/btq054.
- [35] Maes, E., Kelchtermans, P., Bittremieux, W., De Grave, K., et al. "Designing Biomedical Proteomics Experiments: State-of-the-Art and Future Perspectives." In: *Expert Review of Proteomics* 13.5 (Apr. 25, 2016), pp. 495–511. DOI: 10.1586/14789450.2016.1172967.
- [36] Martens, L. "Bringing Proteomics into the Clinic: The Need for the Field to Finally Take Itself Seriously." In: *PROTEOMICS - Clinical Applications* 7 (5-6 June 2013), pp. 388–391. DOI: 10.1002/prca.201300020.

- [37] Martens, L. “Public Proteomics Data: How the Field Has Evolved from Sceptical Inquiry to the Promise of in Silico Proteomics.” In: *EuPA Open Proteomics* 11 (June 2016), pp. 42–44. DOI: 10.1016/j.euprot.2016.02.005.
- [38] Martens, L., Chambers, M., Sturm, M., Kessner, D., et al. “mzML—a Community Standard for Mass Spectrometry Data.” In: *Molecular & Cellular Proteomics* 10.1 (Jan. 1, 2011), R110.000133–R110.000133. DOI: 10.1074/mcp.R110.000133.
- [39] McDonald, W. H., Tabb, D. L., Sadygov, R. G., MacCoss, M. J., et al. “MS1, MS2, and SQT—three Unified, Compact, and Easily Parsed File Formats for the Storage of Shotgun Proteomic Spectra and Identifications.” In: *Rapid Communications in Mass Spectrometry* 18.18 (Sept. 30, 2004), pp. 2162–2168. DOI: 10.1002/rcm.1603.
- [40] McKinney, W. “Data Structures for Statistical Computing in Python.” In: *Proceedings of the 9th Python in Science Conference*. Ed. by van der Walt, S. and Millman, J. Austin, Texas, USA, 2010, pp. 51–56.
- [41] Paulovich, A. G., Billheimer, D., Ham, A.-J. L., Vega-Montoto, L., et al. “Interlaboratory Study Characterizing a Yeast Performance Standard for Benchmarking LC-MS Platform Performance.” In: *Molecular & Cellular Proteomics* 9.2 (Feb. 1, 2010), pp. 242–254. DOI: 10.1074/mcp.M900222-MCP200.
- [42] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., et al. “Scikit-Learn: Machine Learning in Python.” In: *Journal of Machine Learning Research* 12 (Oct 2011), pp. 2825–2830.
- [43] Pichler, P., Mazanek, M., Dusberger, F., Weillböck, L., et al. “SIMPATIQCO: A Server-Based Software Suite Which Facilitates Monitoring the Time Course of LC-MS Performance Metrics on Orbitrap Instruments.” In: *Journal of Proteome Research* 11.11 (Nov. 2, 2012), pp. 5540–5547. DOI: 10.1021/pr300163u.
- [44] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016.
- [45] Rudnick, P. A., Clauser, K. R., Kilpatrick, L. E., Tchekhovskoi, D. V., et al. “Performance Metrics for Liquid Chromatography-Tandem Mass Spectrometry Systems in Proteomics Analyses.” In: *Molecular & Cellular Proteomics* 9.2 (Feb. 1, 2010), pp. 225–241. DOI: 10.1074/mcp.M900223-MCP200.
- [46] Scheltema, R. A. and Mann, M. “SprayQc: A Real-Time LC-MS/MS Quality Monitoring System to Maximize Uptime Using off the Shelf Components.” In: *Journal of Proteome Research* 11.6 (June 1, 2012), pp. 3458–3466. DOI: 10.1021/pr201219e.
- [47] Sharma, V., Eckels, J., Taylor, G. K., Shulman, N. J., et al. “Panorama: A Targeted Proteomics Knowledge Base.” In: *Journal of Proteome Research* 13.9 (Sept. 5, 2014), pp. 4205–4210. DOI: 10.1021/pr5006636.
- [48] Slebos, R. J. C., Wang, X., Wang, X., Zhang, B., et al. “Proteomic Analysis of Colon and Rectal Carcinoma Using Standard and Customized Databases.” In: *Scientific Data* 2 (June 23, 2015), p. 150022. DOI: 10.1038/sdata.2015.22.
- [49] Sturm, M., Bertsch, A., Gröpl, C., Hildebrandt, A., et al. “OpenMS – An Open-Source Software Framework for Mass Spectrometry.” In: *BMC Bioinformatics* 9.1 (Mar. 26, 2008), p. 163. DOI: 10.1186/1471-2105-9-163.
- [50] Sweredoski, M. J., Smith, G. T., Kalli, A., Graham, R. L. J., et al. “LogViewer: A Software Tool to Visualize Quality Control Parameters to Optimize Proteomics Experiments Using Orbitrap and LTQ-FT Mass Spectrometers.” In: *Journal of Biomolecular Techniques* 22.4 (Dec. 2011), pp. 122–126. PMID: 22131886.
- [51] Tabb, D. L. “Quality Assessment for Clinical Proteomics.” In: *Clinical Biochemistry* 46.6 (Apr. 2013), pp. 411–420. DOI: 10.1016/j.clinbiochem.2012.12.003.
- [52] Tabb, D. L., Fernando, C. G., and Chambers, M. C. “MyriMatch: Highly Accurate Tandem Mass Spectral Peptide Identification by Multivariate Hypergeometric Analysis.” In: *Journal of Proteome Research* 6.2 (Feb. 2, 2007), pp. 654–661. DOI: 10.1021/pr0604054.
- [53] Tabb, D. L., Wang, X., Carr, S. A., Clauser, K. R., et al. “Reproducibility of Differential Proteomic Technologies in CPTAC Fractionated Xenografts.” In: *Journal of Proteome Research* 15.3 (Mar. 4, 2016), pp. 691–706. DOI: 10.1021/acs.jproteome.5b00859.
- [54] Taylor, R. M., Dance, J., Taylor, R. J., and Prince, J. T. “Metriculator: Quality Assessment for Mass Spectrometry-Based Proteomics.” In: *Bioinformatics* 29.22 (Nov. 15, 2013), pp. 2948–2949. DOI: 10.1093/bioinformatics/btt510.
- [55] Toby, T. K., Fornelli, L., and Kelleher, N. L. “Progress in Top-down Proteomics and the Analysis of Proteoforms.” In: *Annual Review of Analytical Chemistry* 9.1 (June 12, 2016), pp. 499–519. DOI: 10.1146/annurev-anchem-071015-041550.
- [56] Van der Walt, S., Colbert, S. C., and Varoquaux, G. “The NumPy Array: A Structure for Efficient Numerical Computation.” In: *Computing in Science & Engineering* 13.2 (Mar. 2011), pp. 22–30. DOI: 10.1109/MCSE.2011.37.
- [57] Vizcaíno, J. A., Csordas, A., del-Toro, N., Dianes, J. A., et al. “2016 Update of the PRIDE Database and Its Related Tools.” In: *Nucleic Acids Research* 44 (D1 Jan. 4, 2016), pp. D447–D456. DOI: 10.1093/nar/gkv1145.
- [58] Walzer, M., Pernas, L. E., Nasso, S., Bittremieux, W., et al. “qcML: An Exchange Format for Quality Control Metrics from Mass Spectrometry Experiments.” In: *Molecular & Cellular Proteomics* 13.8 (Aug. 1, 2014), pp. 1905–1913. DOI: 10.1074/mcp.M113.035907.

- [59] Wang, X., Chambers, M. C., Vega-Montoto, L. J., Bunk, D. M., et al. “QC Metrics from CPTAC Raw LC-MS/MS Data Interpreted through Multivariate Statistics.” In: *Analytical Chemistry* 86.5 (Mar. 4, 2014), pp. 2497–2509. DOI: [10.1021/ac4034455](https://doi.org/10.1021/ac4034455).
- [60] Wen, B. and Gatto, L. *proteoQC: An R package for proteomics data quality control*. R package version 1.6.0. 2016.
- [61] Zhang, B., Wang, J., Wang, X., Zhu, J., et al. “Proteogenomic Characterization of Human Colon and Rectal Cancer.” In: *Nature* 513.7518 (July 20, 2014), pp. 382–387. DOI: [10.1038/nature13438](https://doi.org/10.1038/nature13438).