



Data quality-aware genomic data integration

Anna Bernasconi¹

Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milan, Italy

ARTICLE INFO

Keywords:

Data quality
Data integration
Data curation
Genomic datasets
Metadata
Interoperability

ABSTRACT

Genomic data are growing at unprecedented pace, along with new protocols, update policies, formats and guidelines, terminologies and ontologies, which are made available every day by data providers. In this continuously evolving universe, enforcing quality on data and metadata is increasingly critical. While many aspects of data quality are addressed at each individual source, we focus on the need for a systematic approach when data from several sources are integrated, as such integration is an essential aspect for modern genomic data analysis. Data quality must be assessed from many perspectives, including accessibility, currency, representational consistency, specificity, and reliability.

In this article we review relevant literature and, based on the analysis of many datasets and platforms, we report on methods used for guaranteeing data quality while integrating heterogeneous data sources. We explore several real-world cases that are exemplary of more general underlying data quality problems and we illustrate how they can be resolved with a structured method, sensibly applicable also to other biomedical domains. The overviewed methods are implemented in a large framework for the integration of processed genomic data, which is made available to the research community for supporting tertiary data analysis over Next Generation Sequencing datasets, continuously loaded from many open data sources, bringing considerable added value to biological knowledge discovery.

1. Introduction

Genomics is going to generate the largest “big data” problem for the mankind: between 100 million and 2 billion human genomes are expected to be sequenced by 2025 [114]. High-throughput technologies and, more recently, Next Generation Sequencing [111] have brought increasing amounts of genomic data of multiple types, realizing huge steps towards unravelling human genome mechanisms and applying them for unprecedented personalized medicine outcomes. Prior to data analysis and biological knowledge discovery, data and metadata integration is considered an activity of irrefutable priority, with pressing demands for enhanced methodologies of data extraction, matching, normalization, and enrichment, to allow building multiple perspectives over the genome; these can lead to the identification of meaningful relationships, otherwise not perceivable when using incompatible data representations [107].

Bioinformatics, including genomics in particular, operates traditionally by exploiting the considerable fieldwork on data acquisition, wrangling and analysis from its practitioners. Best practices are accumulated across labs and different projects, shared on forums (e.g., <https://www.biostars.org/>, <http://seqanswers.com/>, <https://www.researchgate.net/>) and collected in the documentation

or wiki-guides in code repositories of tools and software. Within these processes, bioinformaticians are mostly concerned with the quality of the experimental data produced by sequencing platforms, for which consolidated pipelines – often composed of many scripts – are available.

In comparison, quality actions that can be performed when aggregating multiple experimental data in systematized ways, have received less attention. However, with the emergence of a culture of data FAIRness [122] and of open and sharable science – promoted by initiatives such as FAIRsharing [109] – caring for data standards in both schemata and values becomes increasingly important, in the same way as implementing integration practices that foster data quality (focusing on accuracy, consistency, currency, and reliability [64]). In its 2012 report on quality of data, IBM found that 1 out of 3 business leaders do not trust the information they use to make decisions (<https://opensistemas.com/en/the-four-vs-of-big-data>); this ratio is unacceptable in fields like health-care and precision medicine, that are strongly driven by genomic databases and decision methods.

Recently, we have observed a trend of initiatives that gather tools and data structures to support interoperability among highly heterogeneous systems, to help bioinformaticians perform a set of curation and annotation operations. These include community-driven efforts

E-mail address: anna.bernasconi@polimi.it

¹ [orcid=0000-0001-8016-5750]

such as [bio.tools](https://www.ebi.ac.uk/bioinformatics/central/services/) [68] (anchored within ELIXIR, <https://www.eelixir-europe.org/>) service providers (EBI [98]), software suites (Bioconductor [67]), or lists (<http://msutils.org/>). Specific instances include APIs such as BioPython [34], tailored scripts, and field descriptions to be parsed (Bioschemas.org [60]). By using, e.g., the EDAM ontology [69], single initiatives can build bridges among resources, while conforming to well-established operations, types/formats of data and application domains.

In this fashion, most problems are handled within a single database by means of *on-the-fly data integration*, driven by a community-inspired guidance. On the other hand, a more systematic approach of *low-level integration* – based upon experience in building solid data warehouses – has also been adopted, helping to reach stable interoperability among imported sources. In these years we witnessed attempts to this kind of approach at many international centers for genomics (including the Broad Institute – <https://www.broadinstitute.org/> – and Wellcome Sanger Institute – <https://www.sanger.ac.uk/> – that are so far unpublished) as well as in companies (including SciDB, implemented by Paradigm4, <https://www.paradigm4.com/technology/scidb-platform-overview/>). In the context of research applied to real problems of the domain, the data-driven Genomic Computing project (GeCo [27]) has dedicated considerable efforts to integrate sources of data that are open for secondary research use, hence downloadable to a common repository, continuously updated. Such systems provide the advantage of offering to users practical work environments. Indeed, biologists and clinicians appreciate ready-to-use repositories, while the know-how of bioinformaticians/developers (on scripting and querying technologies) may not be always at hand.

The emergence of the mentioned positive experiences indicates how data quality can be generally addressed, within thousands of cooperative studies that are jeopardized due to the poor quality of genomic data integration. Poor quality arises at very diverse levels: protocols, data units and dictionaries, metadata models and terminologies. We propose a step forward in the holistic understanding of genomic data and metadata integration, describing a number of methods that can be practically employed to resolve heterogeneity. While most of the introduced issues and techniques have commonalities with general data integration problems, we instantiate them in the specific genomic data context, providing practitioners with easy-to-relate examples to guide their procedures.

In [Section 2](#) we discuss the state of the art since the earliest works on quality-aware genomic databases management [16,91]. In [Section 3](#) we focus on *processed data* (i.e., the signal extracted from raw genomic datasets) and on metadata (i.e., data description), which is the main driver for interoperability and interconnectedness of different databases; in this context, we present a taxonomy of data integration procedures that can positively affect data quality issues. We interpret integration as a set of steps [13], during which practitioners encounter several *heterogeneity loci*, which are contexts that cause heterogeneity and that may be addressed during the specific activities of integration.

In [Section 4](#) we describe a collection of problems with related practical examples and solutions proposed as common practices or specific of our experimented pipelines. In this context, data integration involves: synchronizing the content of a global repository with the data sources, organizing data and corresponding metadata with a unique orthogonal approach, considering interoperability of data descriptions and, more in general, allowing heterogeneous datasets to be used together seamlessly. More pragmatically, we argue that the problem of data quality cannot be addressed as an independent issue. It is entangled with many other aspects regarding data modeling, management, integration and usage. We do not consider quality deriving from original sources as it is not a space where we can intervene *a posteriori*. Instead, we propose a novel angle: addressing data quality dimensions while diverse data sources are being integrated together to enable further applications. In conclusion, in [Section 5](#) we reaffirm the need for quality-aware solutions to integration and mention our vision on upcoming approaches dedicated to this matter.

2. Background

From DNA microarrays [70] to Next Generation Sequencing [111], “quality” in genomics has been usually employed to refer to “quality control” steps on sequences, typically a pre-processing activity aimed at removing adapter sequences, low-quality reads, uncalled bases, contaminants. Instead, in this review we refer to Data Quality (DQ) in the broader sense defined by Wang and Strong [119], usually captured by the expression “fitness for use”, i.e., the ability of datasets to meet their users’ requirements. DQ is evaluated by means of different quality dimensions (i.e., single aspects or components of a data quality concept [115]). State of the art techniques to solve data quality issues in general databases are summarized in [50], under the name of ‘data cleaning’.

In [Figure 1](#) we appreciate the chronological order of publication of relevant literature. In general, more foundational works of data quality in genomics/biological database have appeared in the early years between 2003 and 2008, building the first baseline for this subject, while after 2014 we observe more specific contributions. Müller et al. [91] examine the quality of molecular biological entities databases. Within the production of data, the authors identify intrinsic problems that lead to incorrect data, concluding that traditional data cleaning techniques, used successfully in many other domains, do not fit the peculiarities of genomics. While giving a complete review of potentially very dangerous errors in sequence and annotation genomic databases, the discussion leaves aside processed data as well as aspects related to data integration and integrated access to multiple heterogeneous sources.

While in 2005 Martínez and Hammer propose the conceptual integration of data quality measures inside a model of data [80], the research group led by Berti-Équille is more focused on the problems deriving from warehousing genomic data [16,63,89]. Their overall experience is summarized in [88], where they claim that metadata describing data preparation and data quality are not exploited enough for ensuring valid results of downstream data analysis. Along the same lines of the pioneer work of Wang and Strong [119], some works propose measures and conceptual frameworks that take into account user-driven quality requirements; see the work by Missier et al. [87], BioGuideSRS [35], BioDQ [81] and the more recent paper by Veiga et al. [118]).

A preliminary work by León et al. [76] classifies the data quality properties that are most relevant for genomics; it was then applied concretely to a Crohn’s Disease clinical diagnosis case study [97]. The general framework has been described very recently in Pastor et al. [99]. Rajan et al. [104] have recently proposed to build a knowledge base for assessing quality and characterizing datasets in biomedical repositories, thus including also genomics and other translational research data. Other works address data quality on specific kinds of genomic databases (e.g., by Hedeler and Missier [64] for transcriptomics and proteomics, by Etcheverry et al. [49] for Genome Wide Association Studies, and by Gonçalves and Musen [59] for repositories of biological samples).

As to specific addressed problems, duplicate detection in biological data was dealt with association rule mining first by Koh et al. [74], then by Apiletti et al. [2,3] and by Müller et al. [90], with a focus on contradicting databases. Recent works cover the prevention of redundancy in big data repositories (UniProt KB in [23] and high throughput sequencing in [52]), providing a comparison with other widely used biological large data repositories.

A number of approaches focus on primary data archives quality (i.e., sequence databases such as GenBank [110]) to automatically detect inconsistencies with respect to literature content (see [20,21]) and to provide benchmarks [31], general categorizations of duplicates [32], deduplication clustering methods [30], as well as insights on characteristics, impacts and related solutions to the problem of duplication in biological databases [29].

Finally, data integration in a quality-aware perspective includes practices of data curation [108] and of service/process curation [58]. Data curation is explored in [117], where the Eagle-i system is devel-

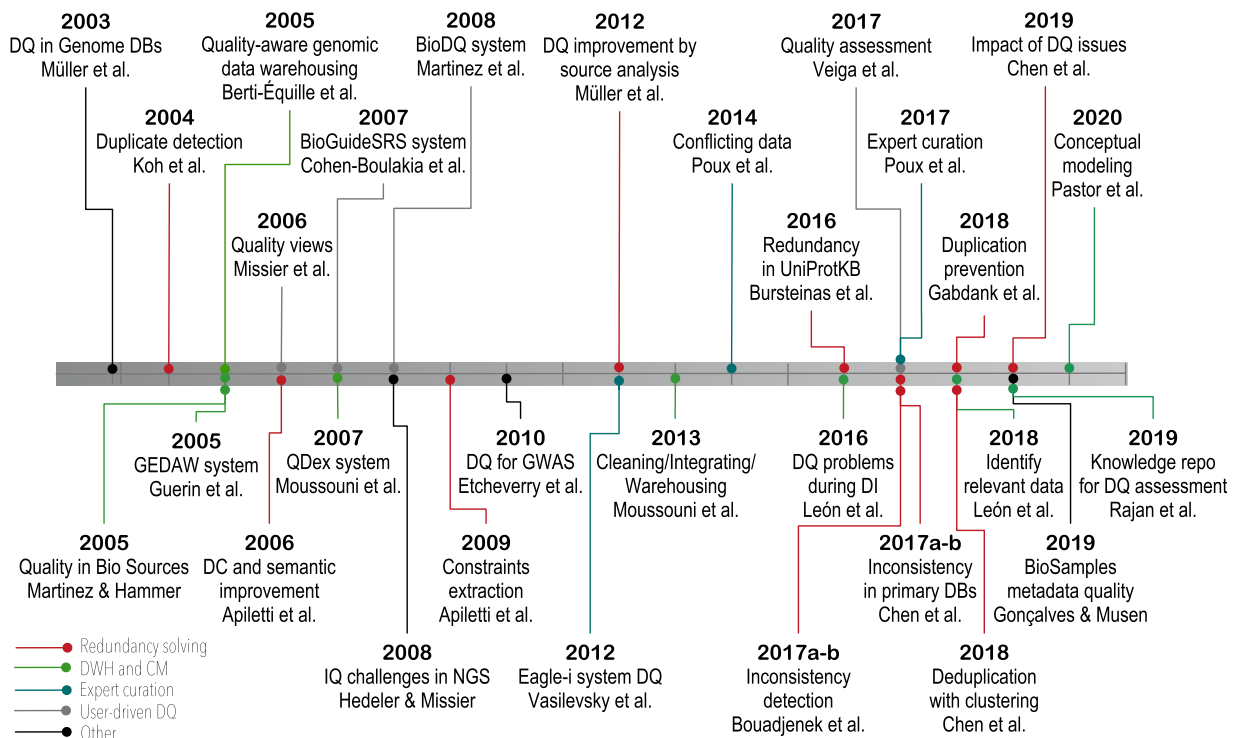


Fig. 1. The timeline of publications targeting data quality issues in biological (and more precisely genomic) databases. Red circles represent works describing approaches to resolve duplication; green circles are works on data warehousing or conceptual modeling; blue circles cover expert curation literature; grey circles are for user-driven data quality approaches; black circles are uncategorized works. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

oped to facilitate collaborative curation, and in [101,102], as a means to deal with conflicting and erroneous data in UniProtKB.

Unfortunately, things have not changed much since the first contributions in this area: data integration is still responsible for solving many data quality problems in this overarching big data challenge. While the focus until now has been much on quality of original data, not much stress has been dedicated to processed data and to metadata issues, which are critical during the data curation process. We thus consider the challenges reported in the mentioned works and remodel them into data quality-driven methods that are already implemented in a working integration framework.

3. Genomics data quality dimensions

The preliminary generation of genomic data follows guidelines and collections of best practices that are gathered during years of practitioners' experience; they are paired by metadata, describing the produced datasets. These are submitted to repositories or collected by consortia that coordinate big research projects and are appointed with the responsibility of publishing it on their platforms. Unfortunately, the integrated use of data coming from different data sources is very challenging, as heterogeneity is met at multiple stages of data extraction (e.g., download protocols, update policies), integration (e.g., conceptual arrangement, values and terminologies), and interlinking (e.g., references and annotation).

While integrating genomic datasets, either for *ad hoc* use in a research study, or for building long-lasting integrated data warehouses, we deal with various complexities that arise during three phases: i) download and retrieval of data from the (potentially multiple) sources; ii) transformation and manipulation, providing fully or partially structured data in machine-readable formats; iii) enrichment, improving the interoperability of datasets.

With *heterogeneity locus* we refer to an activity or phase within the genomic data production/integration process that exhibits heterogeneity

issues, thus undermining the quality of resulting resources. Dividing production from integration, the taxonomy in Fig. 2 keeps track of all the phases in which a genomic data user may need to resolve problems related to non-standardized ways of producing data, making it accessible, organizing it, or enhancing its interoperability. Issues may derive from diverse data and process management habits across different groups that work within the same institution; even more so across different ones. In Fig. 2 the *heterogeneity loci* (listed in the central column) are grouped by production and integration phases (on the left) and are related to data quality dimensions (on the right) that are critical in the represented heterogeneity aspects and are described in the following subsections. In the figure, as in the remainder of the paper, we refer to widely used state-of-the-art definitions of data quality dimensions [105,119] as well as to more recent ones [5,9].

From the broader landscape of processed data sources presented in [14], in this review we focus on example applications that refer to a number of open *general-purpose resources* (Encyclopedia of DNA elements, ENCODE [48]; Roadmap Epigenomics Project [75]; 1000 Genomes Project [33]; Genotype-Tissue Expression Consortium, GTEx [77]; Genome Wide Association Study GWAS Catalog [22]), *cancer genomics resources* (The Cancer Genome Atlas, TCGA [121]; Genomic Data Commons, GDC [62]; International Cancer Genome Consortium, ICGC [126]), *primary archives* (Gene Expression Omnibus, GEO [8]), and *annotation resources* (GENCODE [51] and NCBI RefSeq [95]).

3.1. Accuracy and validity of generated content

Within production, datasets are generated and then published. Generation includes complex practices and challenges, involving quality issues related to *accuracy*, i.e., the degree to which produced experimental data correctly and reliably describe real-world represented events [119]. Such aspects have been thoroughly reviewed in previous works [64,76,91]. Much less investigated, instead, are the issues related to metadata authoring (i.e., preliminary compilation of infor-

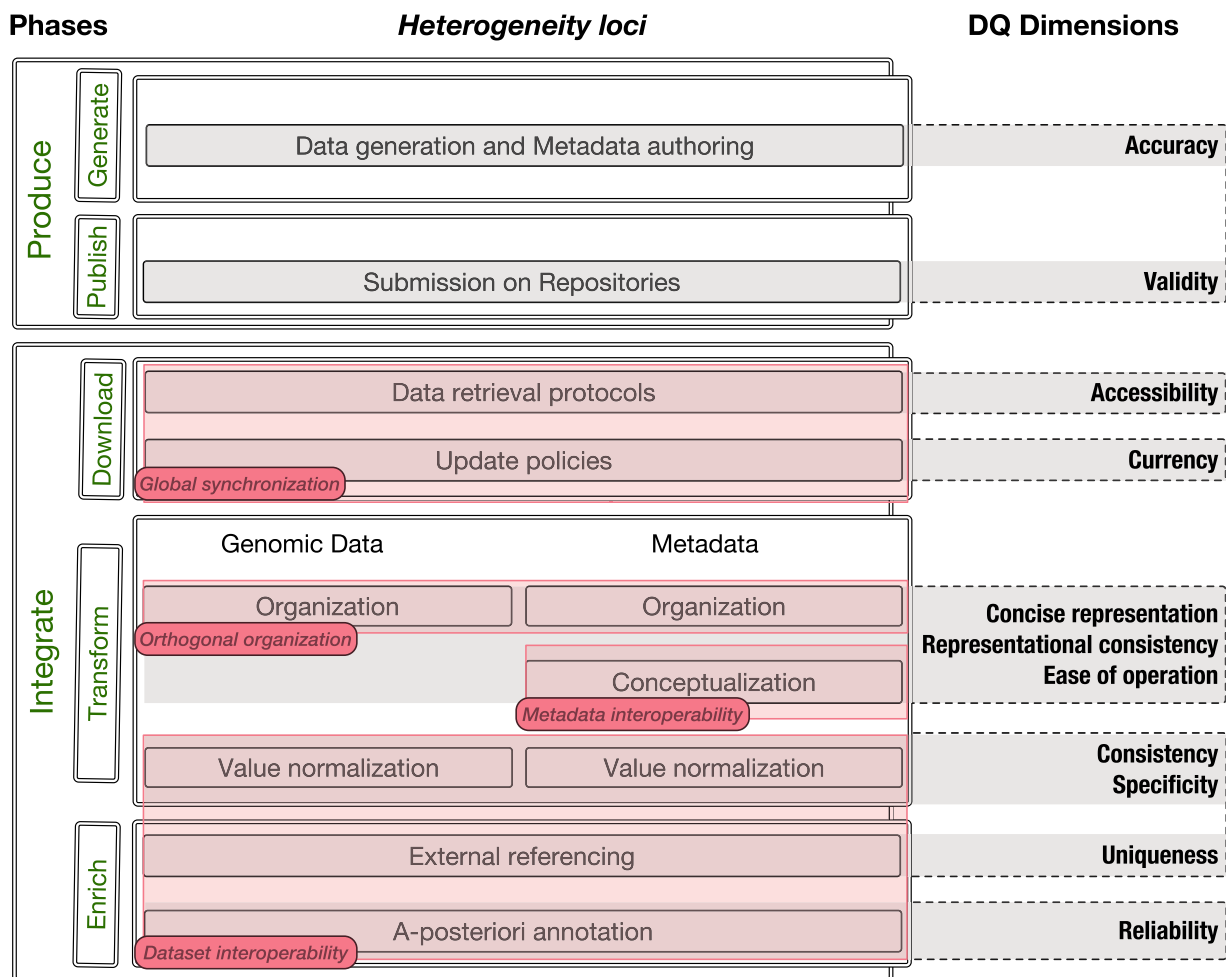


Fig. 2. Taxonomy of *heterogeneity loci* and affected data quality dimensions during genomic data integration; dimensions are equipped with the references where they are primarily defined; they are further discussed in dedicated subsections. Pink rectangles are explained in Sections 4.1–4.4.

mation) [93]. Until very recently, practitioners and investigators from the biomedical community have not recognized metadata creation as a first class activity in their work. As a consequence, *accuracy* of metadata values is negatively affected and it becomes very hard for many final users to work with it.

As publication paves the way to downstream opportunities for integration and analysis, a growing number of scientific journals require, upon submission, that genomic experimental data are contextually submitted to public data repositories [1] (such as GEO, SRA [73] or ArrayExpress [6]). Unfortunately, metadata instances in GEO repository suffer from redundancy, inconsistency, and incompleteness [124], especially due to a lightly regulated submission process. Users are allowed to create arbitrary fields that are not predefined by set dictionaries, many requested information are unstructured, and *validity* of the fields' values (i.e., the degree of their compliance with syntax – format, type, range – of the corresponding definitions [5]) is not checked. Information for submitting high-throughput sequencing data is listed at <https://www.ncbi.nlm.nih.gov/geo/info/seq.html>. A wide literature has been produced to capture structured information from GEO *a posteriori* (e.g., [100,120]). The scenario of alternative repositories, i.e., NCBI BioSample [7] and EBI BioSamples [44], is witnessed in [59].²

² This revealing analysis shows many insights, such as: i) in the description field the concept 'age' appears in 33 different ways (e.g., AGE, Àge, age (yr-old), Age of patient, age_years); ii) 73% of Boolean metadata values are not actually true or false; iii) 26% of integer metadata values cannot be parsed into integers.

Once published on public repositories, data become available for a much wider community, they are potentially re-utilized in secondary analysis or integrated in other platforms; disorganization in the conveyance of provenance information and descriptions of generation procedures negatively affects 'data lineage' [45].

3.2. Accessibility of open genomic data

Sources display diverse download options including programmatic interfaces (APIs), file transfer protocol (FTP) servers, and simple web interface links (HTTP or HTTPS). According to our analysis of important consortia housing open genomic data: i) ENCODE, GDC, and ICGC provide HTTPS API GET/POST services to retrieve lists of files corresponding to chosen filters and additional services to download the corresponding files one by one; ii) Roadmap Epigenomics, GENCODE, RefSeq, 1000 Genomes, and GWAS Catalog store all files on FTP servers, that can be navigated programmatically; iii) GEO provides a variety of methods (both through its own portal and from alternative interfaces), each concentrated on selected partitions of the entire repository content; iv) GTEx can only be accessed from its HTML website.

Only in some cases metadata information is structured and programmatically available. Sometimes metadata files are associated 1:1 to data files (i.e., each data file has a corresponding metadata file); in these cases they can be downloaded in similar ways as the corresponding data file (e.g., by just adding a parameter in an API call, as in ENCODE, or by calling a similar API endpoint using the same file identifier, as in GDC). In other cases, a single metadata file describes a collection of experi-

ments (e.g., Roadmap Epigenomics) or metadata information need to be retrieved in a number of different summary text files, where correspondence between a row and a genomic data file may be obtained using sample IDs (e.g., ICGC or 1000 Genomes). Other times sources have dedicated no effort in systematizing metadata or bringing metadata to a single place; these can only be gathered from descriptions scattered across Web pages.

Accessibility measures the ability of genomic data consumers to easily and quickly retrieve datasets [119]; it is a critical aspect in this phase, as very specific modules need to be created for each source, often upon analysis of cumbersome online documentation and understanding of specific parameters of each portal. Moreover, many well-known open-data databases (such as Cistrome [127], Broad Institute's CCLE [56], and COSMIC [116]) require authentication to access their data; these can only be downloaded and not re-distributed, creating a barrier to integration.

3.3. Currency of retrieved information

Measuring the extent to which data are up to date (the so-called *currency* [105]) is not trivial, as file version synchronization between integrative solutions and original sources strictly depends on the information about the data update state made available in the specific scenarios. The analyzed sources provide such kind of information in different ways: i) ENCODE, GDC and ICGC store information about last data update and checksums within their metadata; ii) Roadmap Epigenomics is a once-for-all project: it will not be updated (at least in the same distribution), therefore does not give such information; iii) 1000 Genomes organizes copies of its data, sequenced in different phases of the project, in different folders of the FTP—update information can be inferred from the paths of the files; iv) GENCODE and RefSeq produce different versions regularly; they are associated to release dates, available in the folders names used on the FTP server; v) GTEx and GWAS Catalog embed the source data version (and subversion) within the file names (e.g., “GTEx_v7_Annotations_SampleAttributesDS.txt” or “gwas_catalog_v1.0_studies_r2019-12-16”).

On the contrary, metadata update information are available only in specific cases. For example let us consider the case of ENCODE and GDC, which have complex hierarchical metadata structures in JSON. ENCODE centers its model on the *Experiment* entity, including *Biosamples* with many *Replicates*, to which *Files* belong. GDC is centered on the *Patient* entity, providing multiple *Samples*; data are also divided by *Project* of a certain *Tumor Type*, for which many *Data Types* are given. These sources associate an update date to each JSON element representing an entity, such as “Experiment”, “Treatment”, “Donor”. The update date automatically pertains also to the elements contained in the entity (e.g., Experiment.assay, Donor.age, Treatment.pipeline...), allowing a fine-grained definition of last update of each single metadata unit. For all sources where files are downloaded from an FTP server, the upload date of files can be used as reference metadata update date.

3.4. Representational conciseness and consistency

Transformation is necessary to organize genomic data and the related descriptions into formats that allow *conciseness* and *consistency* in the *representation* of information. These dimensions respectively measure the ability to compactly, yet completely, represent data and the ability to present data in a same format, allowing backward compatibility [119]. When targeting further data manipulation and analysis, these requirements consequently translate into *ease of operation*, i.e. the extent to which data are easily used and customized [119].

Genomic data organization is a hard task because files have many formats with different semantics (e.g., expression matrices, sets of annotations, sets of peaks measured during an experiment or instead corresponding to a specific reference epigenome...). There does not exist any collectively accepted standard for a general yet basic data unit, that is

able to concisely represent very heterogeneous input data types (given that rows and columns can express different conceptual entities and with different levels of detail).

Also metadata formats are various: hierarchical ones (such as JSON, XML, or equally expressive) adhere to in-house conceptual models; tab-delimited formats (TSV, CSV or Excel/Google Spreadsheets) present different semantics for rows and columns; completely unstructured metadata formats, collected from Web pages or other documentation provided by sources, need to be understood case by case.

3.5. Value consistency, uniqueness and specificity

Heterogeneity is present not only in representation formats, but also evident in values. Normalization activities may involve adding/standardizing genomic coordinates (e.g., from 0-based coordinates to 1-based or vice-versa) and other positional information, adding associated known genomic regions (e.g., genes, transcripts, miRNA) from standard nomenclatures, or formatting into general/source-specific formats, such as narrowPeak or broadPeak ENCODE's standards. A non-exhaustive list of commonly used genomic formats is found at <https://genome.ucsc.edu/FAQ/FAQformat.html>.

Also metadata that describe datasets in different sources are often incompatible or incomplete, using various reference ontologies or no terminology at all. The lack of *consistency* between value domains (i.e., no compliance with semantic rules defined over sets of values [9]) certainly hinders interoperability among sources.

Moreover, as the identity of genomic records is realized using descriptive fields in metadata – usually in addition to internal identifiers – metadata are in charge of handling *uniqueness* with respect to instances within a same source, ensuring that no exact duplicates exist for the same experimental data record [105]. Uniqueness is certainly a goal within single sources, while in the genomics domain (and biomedical more in general), it is accepted that entries representing same real-world entities are repeated in different sources, provided that linking references are present and records are aligned (as debated in [113]). This activity, improving lineage and interoperability of the database content, is very critical especially in an application field where resources are typically not well interlinked and information is only present in some databases and with different degrees of value *specificity* (referred to as *level of detail* in [105]).

3.6. Reliability of annotations

Annotation, i.e. structural and functional classification of (sub)sequences, is an across-the-board activity of the genomic data life cycle. Typically, annotating means associating genomic regions with labels from Gene Ontology [41] (explaining the related molecular function, biological process, cellular component) or with medical concepts related to the sequence (e.g., from UMLS [17]). The process is described in many works [47,63,64], hinting at the related data quality aspects. Annotations are either done by human experts, accurately based on literature evidence and certainly time consuming, or predicted automatically by algorithms that try to infer structural and functional information from similar genes/proteins (worst in terms of accuracy but much less time consuming).

Semantic annotation is instead a typical practice on metadata. As surveyed in Bodenreider [18], ontologies have been widely used in biomedical data management and integration for many years, with the main purpose of improving data interoperability [112]. Many tools are already available to allow semantic annotation with biomedical ontological concepts (see Annotator [71], EBI Zooma (<https://www.ebi.ac.uk/spot/zooma/>), NIH UMLS MetaMap [4], HeTop [61]). Techniques of text-mining [66] have been put into practice on many sources of biomedical text, including abstracts and experiment description from Gene Expression Omnibus [28,53,57], so far one of biggest yet least curated and standardized sources, thus

drawing more attention and efforts. The problem of choosing the right ontologies for semantic enrichment is addressed in [96].

However, guidelines to achieve more standard annotation outcomes are still lacking. *Reliability* [119] of results (i.e., the extent to which annotations can be confidently used to connect and compare datasets) remains a critical aspect of annotation, being dependent on both the algorithm and the acceptance of the ontology in the biomedical community (which itself results from many factors, sometimes hard to measure).

4. Quality-aware methods for data integration

During the research activity documented in [14] we analyzed about 30 data repository hosts, consortia databases, platforms, and interfaces that integrate heterogeneous datasets. We performed various genomic data excavation sessions with the perspective goal of understanding the most important open data sources to be included in a rich processed data repository. Within this process, we experienced several cases of heterogeneity located in the specific *loci* depicted in Fig. 2 (see pink rectangles of different sizes, marked with labels that characterize Sections 4.1-4.4), necessarily resulting into data quality problems. In the following discussion, we focus on the *loci* related to integration phases. For each, first we provide paradigmatic real-world instances. Then, we formalize the problem into overarching questions, specifying the data quality dimensions that are addressed at this stage (as listed in the previous section). Finally, we outline methods that are employed to resolve the issue, from literature and from integration efforts realized in the GeCo project.

4.1. Global repository synchronization with data sources

In the following we provide two example problems regarding data synchronization on the widely employed TCGA and ENCODE sources. Two additional examples, based on ICGC and 1000 Genomes, are available in the Section 1 of Supplementary material.

Example 1. Until 2016, TCGA data was available through a data portal that provided metadata only in XML format, using *biospecimen supplements* and *clinical supplements* that described respectively the biological samples analyzed in the experiments and the patient history, clinical information, and treatments. TCGA has undergone a transition towards the new GDC portal, where the data has been, by now, almost completely transferred. However, there are significant inconsistencies related to metadata. All supplements have been maintained and are still downloadable, but they nowhere fit in the new described data model, available at <https://gdc.cancer.gov/developers/gdc-data-model-0>. Instead, an entirely new collection of metadata, available through programmatic interface, has been defined, divided in four main endpoint groups (*Project*, *Case*, *File*, *Annotation*). The documentation of available fields is at https://docs.gdc.cancer.gov/API/Users_Guide/Appendix_A_Available_Fields/. GDC migration is still ongoing; nevertheless documentation is not consistently updated and it is common to find fields that are already visible in the interface facets (and APIs) but that indeed have null values for all instances in the database. Moreover, not all datasets that were available in the previous portal are now available in the new portal. For these reasons, synchronizing the content of an integrated repository with the one of GDC becomes very critical.

Example 2. ENCODE source elements in JSON schemata, used for searching metadata through Elasticsearch (<https://www.elastic.co/>) are changed very often, as documented in about 90 Changelogs, one for each JSON entity corresponding to a profile. A complete list of ENCODE's data model entities (i.e., profiles) is at <https://www.encodeproject.org/profiles/>. However, metadata instances change also their values. For example, to keep track of the change of about 10 attribute-value pairs in the experiment ENC635OSG (<https://www.encodeproject.org/experiments/ENC635OSG/>), only a simple comment in the metadata was added (i.e., *Submitter*

comment: "IMPORTANT! Bioreplicate 2 was previously annotated as liver from a 4 year old female. It has now been corrected to be liver from a 32 year old adult male.").

Problem formulation. *How can changes on genomic data sources be taken into account to be reflected on integrated repositories, guaranteeing 'currency'? How can it be done systematically, overcoming 'accessibility' issues?*

Method 1 – Source partitioning. When targeting integrated systems up to date, the main difficulty is to identify data partitioning schemes specific for each source (as discussed in [13]); a partition can be repeatedly accessed and source files that are modified within the partition (or added to it) can be recognized, avoiding selectively the download of the source files that are not changed. Suppose we are interested in downloading a certain updated ENCODE portion (e.g., transcriptomics experiments on human tissue, aligned to reference genome hg19). We produce an API request to the endpoint <https://www.encodeproject.org/matrix/?type=Experiment>, specifying the parameters `type = Experiment`, `replicates.library.biosample.donor.organism.scientific_name = Homo+sapiens`, `status = released`, `assembly = hg19`, and `assay.slims = Transcription`. In 1000 Genomes, as there are no API available, we instead navigate the FTP server directly and check the most updated release available on ftp://1000genomes.ebi.ac.uk/vol1/ftp/data_collections/1000_genomes_project/release; we consequently enter the relative folder (e.g., `20190312_biallelic_SNV_and_INDEL`) and download all chromosomes files.

Method 2 – Event-based update. We periodically check source websites and FTP servers for new data. We use a relational database (called `importer_db` in the following) to manage the synchronization process between the data sources and our local repository. The `Sources` table has many `Datasets`, each of which corresponds to `Files` (i.e., the genomic region data files). Each `Run` of the download process checks unique properties of data `Files` such as `URL`, `Origin-LastUpdate`, `OriginSize` and `Hash`, used to compare the local copy of the file with the original one on the data source: new files are stored and processed; missing files (i.e., deprecated on the source) are copied to an archive; matching files that have identifying values different from the corresponding local values stored in the `importer_db` are re-downloaded.

For different sources *ad hoc* software modules can be developed to periodically check for changes in the Changelog, schema definitions, documentation, in search for motivation to update our local copy of the source data. As an example, GEO offers to registered users a mechanism to be notified when new data is available, relevant for a search saved previously (<https://www.ncbi.nlm.nih.gov/geo/info/faq.html#notifications>). TCGA2BED [46] was realized to handle data acquisition and transformation for TCGA source and OpenGDC [26] provides an updated framework to acquire synchronized data also from GDC portal.

4.2. Orthogonal data and metadata organization

Examples. While there is common agreement on the terminology used to define genomic data types (e.g., mutation, copy number variation, chromatin accessibility), data types are typically not rendered using the same machine-readable formats (e.g., there exist both VCF-like and ICGC-like mutation formats (<https://docs.icgc.org/submission/guide/icgc-simple-somatic-mutation-format/>), gene expression data may be presented as sample/gene matrices or just as lists of genes with expression values per aliquot). Sometimes formats are defined at "experiment time" to suit particular needs of the data; they are documented in plain text attachments. An example of format defini-

tion of ENCODE tsv files representing gene expression matrix is <https://www.encodeproject.org/documents/c2bbc04-9b9d-41aa-883f-bbba9bc45e68/>. In this kind of documents, some specifications may further confuse data organization, as matrix cell values are allowed to contain *ad hoc* formatting semantics (e.g., from an ENCODE format definition document: “The value in the cell contains two strings, one for TPM values and another for FPKM values, separated by underscore; each string contains values for each replicate, separated by colon.”).

Many formats also fail at keeping representation levels orthogonal; for example, properties that represent values aggregated over a multitude of regions are sometimes displayed as part of single regions, repeated in each of them. In 1000 Genomes variation data, each line expresses one mutation and contains, as a property, the measure of allele frequencies across entire geographic populations (i.e., thousands of samples).

Additionally, data from specific projects are simultaneously provided by different portals, that however re-shape it in several ways: ENCODE portal includes Roadmap Epigenomics data, re-processed using distinct pipelines and with completely different data schemata and metadata; TCGA data appears in both GDC and ICGC with very dissimilar representation both for data (one textual file for each aliquot from a patient, as opposed to one big spreadsheet containing independent lines, each connected to a patient) and for metadata.

Problem formulation. There is no agreement towards a basic genomic data unit for tertiary analysis. A common choice is to prepare one file for each experimental session; lines are genomic regions associated to some properties. Other times data units are huge matrices of patients or samples crossed with genes, miRNA, or other encoded sequences. Each source and each data type, thus, needs its own “basic unit”. *Can genomic data be expressed using a unique model that is general enough to represent all analyzed formats (‘concise and consistent representation’), and that also allows ‘ease of operation’?*

Method – Genomic Data Model and sample identity. The need for defining a *genomic basic data unit* is emerging: a single piece of information that contains genomic regions with their properties and is identifiable with an entity that is interesting for downstream analysis (e.g., a patient, a biological sample, a reference epigenome...). Any set of downloaded files – with their input format – should be convertible through a *transformation relation* into a set of genomic basic data units. We define as *transformation relation cardinality* the pair $X : Y$, where X is the cardinality of the set of files from the input source and Y is the cardinality of the output set of “basic units” into which the input is transformed for downstream use in an integrative system; $X : Y$ is a fraction in lowest terms.

As a paradigm that more generally includes the interval-based genomic data representations (see BEDTools [103] and BEDOPS [94] for example), an interesting candidate for expressing such basic unit is represented by the *sample* of the Genomic Data Model (GDM, [84]). A sample can express heterogeneous DNA features, such as variations (e.g., a mutation in a given DNA position), peaks of binding or expression (i.e., genomic regions with higher read density), or structural properties of the DNA (e.g., break points, where the DNA is damaged, or junctions, where the DNA creates loops). GDM is based on the notion of *dataset*, i.e. a collection of *samples*. A sample, in turn, consists of two parts: the *region data*, describing the characteristics and DNA location of genomic features, and the *metadata*, describing general properties of the sample, in the form of key-value pairs; in GDM format there is one metadata file for each region data file.

Some sources provide a data file for each experimental event, for example ENCODE. In this case, the transformation has a 1:1 cardinality, i.e., to each ENCODE produced file, it corresponds one GDM sample. Other sources include more complex formats, such as MAF, VCF, and gene expression matrices. In these cases, the transformation phase takes care of compiling one single data file for each patient or univocally identified sample in the origin data. The transformation cardinality is thus 1:N, N being the number of patients or biological samples.

Metadata also feature diverse formats in the analyzed sources: i) hierarchical formats (JSON, XML, or equally expressive) require applying a flattening procedure to create key-value pairs—the key results from the concatenation of all JSON/XML elements from the root to the element corresponding to a value; ii) tab-delimited formats (TSV, CSV or Excel/Google Spreadsheets) strictly depend on the semantics of rows and columns (e.g., 1 row = 1 epigenome, 1 row = 1 biological sample)—they often require pivoting tab-delimited columns into rows (which corresponds to creating key-value pairs); iii) two-columns tab-delimited formats (such as GEO’s SOFT files) are translated into GDM straightforwardly; iv) completely unstructured metadata formats, collected from Web pages or other documentation provided by sources, need case-specific manual processing.

Table 1 shows transformation relation cardinalities regarding both data and metadata input formats, targeting the GDM output format. We analyzed different data types in a number of important data sources, that possibly include files with different formats. Note that, while for descriptive purposes we indicate physical formats (e.g., TSV, TXT, JSON), the indication of cardinalities also embeds a semantic information: how many data units are represented in one file. Following the mapping from input sources into GDM format we can solve systematically the heterogeneity of data formats and prepare the GDM datasets as sets of GDM samples that are uniform in their schema. The Supplementary material (Section 2) provides additional details on this method.

4.3. Metadata interoperability

Examples. Metadata heterogeneity can also be analyzed from other perspectives. From a schema point of view (i.e., how each piece of information is identified and interrelated with others), when searching for disease-related attributes, we find diverse possibilities: “Disease type” in GDC, “Characteristics–tissue” in GEO, “Health status” in ENCODE. From the values point of view, when searching for breast cancer-related information, we find multiple expressions, pointing to comparable samples, e.g., “Breast Invasive Carcinoma” (GDC), “breast cancer ductal carcinoma” (GEO), “Breast cancer (adenocarcinoma)” (ENCODE).

Roadmap Epigenomics expresses ages of samples using a unique column “AGE (Post Birth in YEARS/Fetal in GESTATIONAL WEEKS/CELL LINE CL)” together for three different kind of classes and, consequently, with three different measure units. Example values for single instances are “Fetus (GW unknown)”, “CL”, “Unknown, Unknown, 45Y”, or “49Y, 59Y, 41Y, 25Y, 81Y”, where 5 values are put together to express that the related epigenome is derived from 5 individuals. Also the information about donors and its interrelation with other attributes is confusing. The column “Single Donor (SD) / Composite (C)” discriminates between epigenomes deriving from one or more donors. Yet, in the column “DONOR / SAMPLE ALIAS” (containing identifiers), it happens that an epigenome labeled as deriving from a single donor, contains instead multiple IDs. No further explanation is available to clarify the semantics; other dependent columns such as “sex” and “ethnicity” become also unclear. Paradoxically, one donor, identified by the string H-22772, turns out to be present in two different epigenomes, one derived from lung and one from heart tissue. This problem is easily propagated to other sources, as Roadmap Epigenomics experiments are replicated in ENCODE repository. Here, different experiments present the same external tags (e.g., `roadmap-epigenomics:UW H22772`).

Problem formulation. *How can metadata be employed to provide a global and integrated view of data sources (‘representational consistency’)? Does such representation help users in querying data straightforwardly (‘ease of operation’)?*

Method – Genomic Conceptual Model for metadata normalization. In literature there are works that use conceptual modeling to better explain relations between biological entities [63,97,106]. However, conceptual modeling can serve brilliantly also the purpose of organizing metadata from heterogeneous sources into one global view. The Genomic Conceptual Model (GCM, [15]) is an Entity-Relationship model

Table 1

Census of 13 important data sources reporting for each: the processed data types that can be downloaded (along with meta-data), their physical formats, and the semantic cardinality of the transformation relation with respect to the GDM output format [84].

Source	Data type	Data format, cardinality ^a	Metadata format, cardinality ^a
ENCODE	peaks	BED, 1:1	JSON, 1(experiment):#samples
	transcription	TSV, 1:1	JSON, 1(experiment):#samples
	transcription	GTF, 1:1	JSON, 1(experiment):#samples
GDC	mutations	MAF, 4:1	JSON, 4:1
	gene expression	TXT, 3:1	JSON, 3:1
	methylation, cnv, quantifications	TXT, 1:1	JSON, 1:1
ICGC	mutation, methylation, miRNA/gene expression	TSV, 1:#donors	TSV, 1:#donors
Roadmap Epigenomics	peaks	BED, 1:1	Spreadsheet, 1:(#samples×#epigenomes) ^b
	transcription	TSV, 2:#epigenomes	Spreadsheet, 1:(#samples×#epigenomes) ^b
GENCODE	annotations	GTF, 1:#annotation_types	region data file + webpage, X:1 ^c
RefSeq	annotations	GFF, 1:#annotation_types	region data file + webpage, X:1 ^c
1000 Genomes	variation	VCF, 23:#individuals	TSV, 4:#individuals
GEO	expression	BED, 1:1	HTML/SOFT, 1:#files_from_sample
GTEX	expression	GCT, 1:#donors	TXT, 1:#donors
GWAS Catalog	associations	TSV, 1:1	-
CISTROME	peaks	BED, 1:1	TSV, 1:1
CLE	various	GCT/TXT, 1:#cell_lines	TXT, 1:#cell_lines
COSMIC	various	TSV, 1:#individuals	TSV, X:1 ^c

^a Expressed as $X : Y$; this ratio represents the number X of data (resp. metadata) units used in the origin source to compose Y data (resp. metadata) file(s) in GDM format.

^b Each reference epigenome is used for many data types, thus many GDM samples. The same epigenome-related metadata is replicated into many samples.

^c In these cases it is difficult to build a numerical relation—many meta are retrieved from the data files themselves, in addition to manually curated information.

used to describe metadata of genomic data sources. The main objective of GCM is to recognize a common set of concepts (about 40) that are semantically supported by most genomic data sources, although with very different syntax and forms. GCM is a star-schema – inspired to classic data marts [19] – centered around the ITEM entity, representing a genomic basic data unit, such as the GDM elementary sample. The four dimensions of the star describe the *biology* of the experiment, the used *technology*, its *management* aspects, and the *extraction* parameters for internal organization of items. A complete integration framework (described in [13]) can be employed to download, transform, clean and integrate metadata at the schema level, importing them into the relational database `gcm_db` that implements the GCM physically. Data constraints checks (name existence and value dependencies in [15]) are performed based on a set of manually introduced rules, but also on automatically generated ones, inspired by the works on data cleaning using association rules mining [3] and much in the fashion of [82], who uses rules as a means to generate recommendations for suitable metadata additions to datasets. The conceptual representation of GCM widely helped domain users in finding data more easily from a unique query interface, without having to deal with heterogeneous access points, metadata formats and models, as demonstrated in [10].

4.4. Large-scale dataset interoperability

Example on data. Within ICGC gene annotation is not consistent among different datatypes (e.g., sequence-based gene expression datasets use Ensembl Gene IDs [125], like ENCODE gene quantification data and TCGA gene expression quantification, while array-based gene expression datasets use the gene name convention of HGNC [123]). Within annotation databases themselves, data may be incomplete. For example in GENCODE's comprehensive gene annotation files (ftp://ftp.ebi.ac.uk/pub/databases/genocode/Gencode_human/) not all exons and transcripts regions have a corresponding gene region that includes them. While searching for correct coordinates of a gene, users may alternatively calculate the *start* as the one of its left-most transcript/exon and the *stop* as the one of its right-most transcript/exon,

but this procedure could be not always accurate. Such shortcomings are consequently propagated in all processed data sources where the reference gene annotation is used to codify signals data (e.g., ENCODE). Furthermore, secondary sources use different releases versions to annotate different files (to date, GENCODE has 34 releases, out of which only 6 are still maintained for the new GRCh38 assembly). This makes it hard to consistently compare files from a same source that have been annotated using different reference sets.

Example on metadata. Metadata are affected by the even more complicated issue of ontology misalignment. Ontology CL [86] and EFO [79] reference same concepts: the specific instances in the two ontologies have differences in the values and schema. NCIT [42] and UBERON [92], both including parts of the human body, also show inconsistencies: while “hypothalamus” is considered a synonym of “BRAIN” in NCIT, it is a sub-concept of “brain” in UBERON (five-levels more specific, traversing both relationships of subsumption `is_a` and containment `part_of`).

Using ontologies as a base for further semantic annotation, many algorithms still produce a relevant number of inaccurate annotations (see [28,54]), which result in harder work for the downstream integration process.

Problem formulation. *How can datasets understand each other? Can we normalize data with respect to commonly adopted terminologies ('consistency', 'specificity' and 'uniqueness') and confidently exploit the currently available external resources ('reliability')?*

Method 1 – Data enrichment. A fruitful approach with annotation is the inclusive one: integrators may add as many information as possible, considering the most accepted resources in the field. For structural and functional annotation of genomic regions and sequences, including adding for example gene/transcript/exon identifiers or biological process related to a protein, multiple reference databases may be queried (RefSeq, GENCODE, Ensembl, Entrez [78], HGNC), as documented in TCGA2BED [46] and OpenGDC [26], or performed during the integration process of several datasets from Roadmap Epigenomics and transcriptomics data from ENCODE. Large-scale data integration in genomics can be achieved using cross-references (see [55]); its success

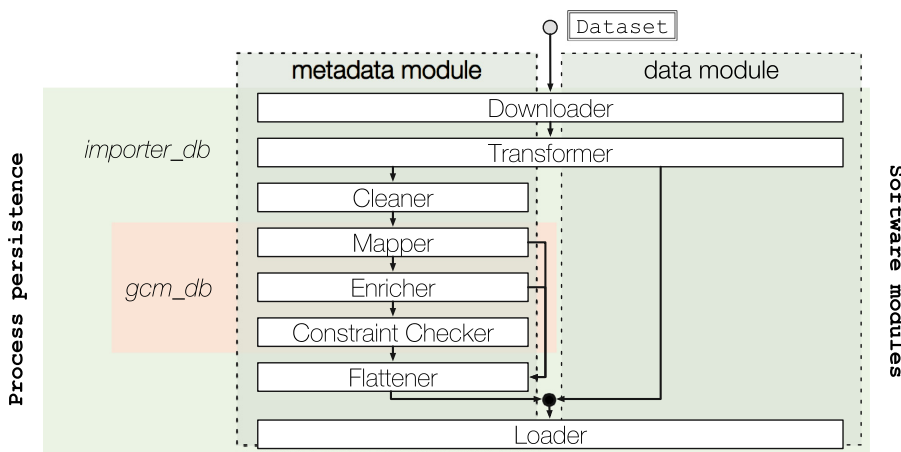


Fig. 3. GeCo integration process, from the *download* of a source partition (based on the prior definition of a GDM dataset), to its *transformation* and all following phases. *Cleaning, mapping, enriching* and *checking* are only performed on the metadata, while data are left unchanged. Metadata are *flattened* from the GCM relational implementation back to the file-based GeCo repository. Data and corresponding metadata of each dataset are *loaded* into a cloud-based data engine for queries on genomic region and metadata [83].

strictly depends on a correct use of persistent identifiers [85]. See the Supplementary material Section 3 for more details on this method.

Method 2 – Metadata enrichment. The process of annotating existing structured metadata with ontological terms, their definitions, synonyms, ancestors, and descendants can be done in an iterative way, automated with respect to the querying of online annotation systems and semantic match computation, but also assisted by an expert manually checking the obtained links [12]. This enhancement of metadata (using specialized biomedical ontologies) can be seen as the construction of a knowledge graph of the content of the repository [11]; it is useful to instrument the search of datasets described by such metadata in a semantically enriched fashion (see GenoSurf interface [25]). See the Supplementary material Section 4 for more details on this method.

5. Discussion and outlook

The integration of data and metadata is of growing relevance in biomedical fields (including genomics), because critical decisions in the domains of health-care – such as precision medicine – depend on it. As individualized predictions become more difficult, they require approaches that combine multiple sources and multiple data types (from genomics, transcriptomics, epigenomics, etc.), possibly completed with clinical data. Heterogeneity aspects affect many actors and stages of the data life cycle. In such situations, data quality dimensions can adequately lead the analysis of problems and related solutions. We have reviewed works that have contributed to data quality-driven approaches in genomics; even with community-driven approaches that propose on-the-fly data integration, the focus has so far been on quality of origin data sources and not so much on the overall process that channels data together for subsequent use. Thus, we have introduced a novel perspective: we have shown a taxonomy of integration phases that directly impact quality of genomic databases and interfaces during data integration; we have detailed the issues related to such phases, providing examples, questions to be addressed, and methods that we experimented during the creation of a repository of high quality, which inspired the discussions of this review paper.

The repository, currently with more than 250k processed items, results from the GeCo project effort. Fig. 3 shows the sequential software modules (<https://github.com/DEIB-GEKO/Metadata-Manager/>) to integrate genomic sources, by solving all the analyzed heterogeneity aspects. Phases are recorded in the `importer_db`: a given dataset is downloaded and periodically synchronized with the origin source, transformed into the GDM format (achieving orthogonal data/metadata organization); metadata are cleaned, simplifying redundant attribute names, mapped into the `gcm_db` (unique conceptual representation, towards interoperability of metadata), semantically enriched and checked

with respect to constraints. The relational representation is flattened to load datasets into a file-based engine for further biological querying [83].

While the described approaches have been successfully implemented in practical contexts [13], future challenges include applying the proposed solutions to complex contexts such as the one of clinical data and translational medicine, that ultimately will need to be also iterated with genomic data. We are aware of important work that is being conducted in parallel on health data [24,38,40,65,72], also employing the data warehouse paradigm as a guarantor of up-to-date de-duplicated data within a public network of research centers [37], usually oriented to support analytics [39]. Several works already address data quality for precision medicine [36,43,97], reviewing the use of genomic data in the medical context, whereas my review is focused primarily on issues of quality in genomic data integration (data comparability, metadata definitions, data standards,...) encompassing all possible uses of genomic data.

In this review, we have shown how resolving quality while building a data repository can effectively create usable integrated environments for researchers. Since many of the described approaches may be useful for other researchers – even in dynamic data integration assets – these will be provided through convenient external programmatic access. Starting from this baseline, we envision a data integration process that includes seamlessly evaluation of quality parameters, towards data and information that are more directly employable in genomic analysis and biological discovery. Predictably, future data integration approaches will include more and more a data quality-aware *modus operandi* with the following characteristics: i) currency-driven synchronization of sources, ii) concise/orthogonal/common data representations, iii) light and interoperable data descriptions, iv) reliability-tailored dataset linkage. All in all, this review highlights trends in genomic data and information integration, which will ideally guide and improve future efforts and activities.

Funding

This work was supported by the European Research Council Executive Agency under the EU Framework Programme Horizon 2020, ERC Advanced Grant number 693174 GeCo (data-driven Genomic Computing).

Declaration of Competing Interest

The author declares no conflict of interest.

Acknowledgments

The author would like to thank Professor Stefano Ceri and Professor Cinzia Cappiello for fruitful discussions and for providing precious suggestions and inspiration during the preparation of the manuscript.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.cmpbup.2021.100009](https://doi.org/10.1016/j.cmpbup.2021.100009)

References

- [1] Microarray standards at last, *Nature* 419 (6905) (2002) 323, doi:10.1038/419323a.
- [2] D. Apiletti, G. Bruno, E. Ficarra, E. Baralis, Data cleaning and semantic improvement in biological databases, *J. Integr. Bioinform.* 3 (2) (2006) 219–229.
- [3] D. Apiletti, G. Bruno, E. Ficarra, E. Baralis, Extraction of constraints from biological data, in: *Biomedical Data and Applications*, Springer, 2009, pp. 169–186.
- [4] A.R. Aronson, F.-M. Lang, An overview of metapap: historical perspective and recent advances, *J. Am. Med. Inform. Assoc.* 17 (3) (2010) 229–236.
- [5] N. Askham, D. Cook, M. Doyle, H. Fereday, M. Gibson, U. Landbeck, R. Lee, C. Maynard, G. Palmer, J. Schwarzenbach, The Six Primary Dimensions for Data Quality Assessment, DAMA UK Working Group, 2013. Online (Jan 12th, 2021, date last accessed), URL: https://www.whitepapers.em360tech.com/wp-content/files_mf/1407250286DAMAUKDQDimensionsWhitePaperR37.pdf
- [6] A. Athar, A. Füllgrabe, N. George, H. Iqbal, L. Huerta, A. Ali, C. Snow, N.A. Fonseca, R. Petryszak, I. Papatheodorou, et al., ArrayExpress update—from bulk to single-cell expression data, *Nucl. Acids Res.* 47 (D1) (2018) D711–D715.
- [7] T. Barrett, K. Clark, R. Gevorgyan, V. Gorelenkov, E. Gribov, I. Karsch-Mizrachi, M. Kimelman, K.D. Pruitt, S. Resenchuk, T. Tatusova, et al., Bioproject and biosample databases at ncbi: facilitating capture and organization of metadata, *Nucl. Acids Res.* 40 (D1) (2011) D57–D63.
- [8] T. Barrett, S.E. Wilhite, P. Ledoux, C. Evangelista, I.F. Kim, M. Tomashevsky, K.A. Marshall, K.H. Phillippy, P.M. Sherman, M. Holko, et al., Ncbi geo: archive for functional genomics data sets—update, *Nucl. Acids Res.* 41 (D1) (2012) D991–D995.
- [9] C. Batini, M. Scannapieco, Data and Information Quality: Dimensions, Principles and Techniques, first ed., Springer Publishing Company, Incorporated, 2016.
- [10] A. Bernasconi, A. Canakoglu, S. Ceri, Exploiting conceptual modeling for searching genomic metadata: a quantitative and qualitative empirical study, in: G. Guizzardi, F. Gailly, R. Suzana Pitangueira Maciel (Eds.), *Advances in Conceptual Modeling*, Springer International Publishing, Cham, 2019, pp. 83–94.
- [11] A. Bernasconi, A. Canakoglu, S. Ceri, From a conceptual model to a knowledge graph for genomic datasets, in: A.H.F. Laender, B. Pernici, E.-P. Lim, J.P.M. de Oliveira (Eds.), *Conceptual Modeling*, Springer International Publishing, Cham, 2019, pp. 352–360.
- [12] A. Bernasconi, A. Canakoglu, A. Colombo, S. Ceri, Ontology-driven metadata enrichment for genomic datasets, in: C.J.O. Baker, A. Waagmeester, A. Splendiani, O.D. Beyan, M.S. Marshall (Eds.), *International Conference on Semantic Web Applications and Tools for Life Sciences*, CEUR Workshop Proceedings, 2275, 2018.
- [13] A. Bernasconi, A. Canakoglu, M. Masseroli, S. Ceri, META-BASE: a novel architecture for large-scale genomic metadata integration, *IEEE/ACM Trans. Comput. Biol. Bioinform.* (2020), doi:10.1109/TCBB.2020.2998954.
- [14] A. Bernasconi, A. Canakoglu, M. Masseroli, S. Ceri, The road towards data integration in human genomics: players, steps and interactions, *Brief. Bioinform.* 22 (1) (2021) 30–44, doi:10.1093/bib/bbaa080.
- [15] A. Bernasconi, S. Ceri, A. Campi, M. Masseroli, Conceptual modeling for genomics: building an integrated repository of open data, in: H.C. Mayr, G. Guizzardi, H. Ma, O. Pastor (Eds.), *Conceptual Modeling*, Springer International Publishing, Cham, 2019, pp. 325–339.
- [16] L. Berti-Équille, F. Moussouni, Quality-aware integration and warehousing of genomic data, in: *Proceedings of the 2005 International Conference on Information Quality (MIT ICIQ Conference)*, Sponsored by Lockheed Martin, MIT, Cambridge, MA, USA, November 10–12, 2006, 2005.
- [17] O. Bodenreider, The unified medical language system (umls): integrating biomedical terminology, *Nucl. Acids Res.* 32 (suppl_1) (2004) D267–D270.
- [18] O. Bodenreider, Biomedical ontologies in action: role in knowledge management, data integration and decision support, *Yearb. Med. Inform.* (2008) 67.
- [19] A. Bonifati, F. Cattaneo, S. Ceri, A. Fuggetta, S. Paraboschi, Designing data marts for data warehouses, *ACM Trans. Softw. Eng. Methodol.* 10 (4) (2001) 452–483.
- [20] M.R. Bouadjene, K. Verspoor, J. Zobel, Automated detection of records in biological sequence databases that are inconsistent with the literature, *J. Biomed. Inform.* 71 (2017) 229–240.
- [21] M.R. Bouadjene, K. Verspoor, J. Zobel, Literature consistency of bioinformatics sequence databases is effective for assessing record quality, *Database* (2017).
- [22] A. Buniello, J.A.L. MacArthur, M. Cerezo, L.W. Harris, J. Hayhurst, C. Malangone, A. McMahon, J. Morales, E. Mountjoy, E. Sollis, et al., The NHGRI-EBI GWAS Catalog of published genome-wide association studies, targeted arrays and summary statistics 2019, *Nucl. Acids Res.* 47 (D1) (2018) D1005–D1012.
- [23] B. Bursteinas, R. Britto, B. Bely, A. Auchincloss, C. Rivoire, N. Redaschi, C. O'Donovan, M.J. Martin, Minimizing proteome redundancy in the uniprot knowledgebase, *Database* (2016).
- [24] T.J. Callahan, A.E. Bauck, D. Bertoch, J. Brown, R. Khare, P.B. Ryan, J. Staab, M.N. Zozus, M.G. Kahn, A comparison of data quality assessment checks in six data sharing networks, *eGEMs (Generating Evidence & Methods to improve patient outcomes)* 5 (1) (2017).
- [25] A. Canakoglu, A. Bernasconi, A. Colombo, M. Masseroli, S. Ceri, GenoSurf: meta-data driven semantic search system for integrated genomic datasets, *Database* (2019).
- [26] E. Cappelli, F. Cumbo, A. Bernasconi, A. Canakoglu, S. Ceri, M. Masseroli, E. Weitschek, OpenGDC: unifying, modeling, integrating cancer genomic data and clinical metadata, *Appl. Sci.* 10 (18) (2020) 6367.
- [27] S. Ceri, A. Bernasconi, A. Canakoglu, A. Gulino, A. Kaitoua, M. Masseroli, L. Nanni, P. Pinoli, Overview of GeCo: a project for exploring and integrating signals from the genome, in: *International Conference on Data Analytics and Management in Data Intensive Domains*, Springer, 2017, pp. 46–57.
- [28] G. Chen, J.C. Ramirez, N. Deng, X. Qiu, C. Wu, W.J. Zheng, H. Wu, Restructured geo: restructuring gene expression omnibus metadata for genome dynamics analysis, *Database* (2019).
- [29] Q. Chen, R. Britto, I. Erill, C.J. Jeffery, A. Liberzon, M. Magrane, J.-i. Onami, M. Robinson-Rechavi, J. Sponarova, J. Zobel, et al., Quality matters: Biocuration experts on the impact of duplication and other data quality issues in biological databases, *bioRxiv* (2019b) 788034.
- [30] Q. Chen, Y. Wan, X. Zhang, Y. Lei, J. Zobel, K. Verspoor, Comparative analysis of sequence clustering methods for deduplication of biological databases, *J. Data Inf. Qual.* 9 (3) (2018) 1–27.
- [31] Q. Chen, J. Zobel, K. Verspoor, Benchmarks for measurement of duplicate detection methods in nucleotide databases, *Database* (2017).
- [32] Q. Chen, J. Zobel, K. Verspoor, Duplicates, redundancies and inconsistencies in the primary nucleotide databases: a descriptive study, *Database* (2017).
- [33] L. Clarke, X. Zheng-Bradley, R. Smith, E. Kulesha, C. Xiao, I. Toneva, B. Vaughan, D. Preuss, R. Leinonen, M. Shumway, et al., The 1000 genomes project: data management and community access, *Nat. Methods* 9 (5) (2012) 459.
- [34] P.J. Cock, T. Antao, J.T. Chang, B.A. Chapman, C.J. Cox, A. Dalke, I. Friedberg, T. Hamelryck, F. Kauff, B. Wilczynski, et al., Biopython: freely available python tools for computational molecular biology and bioinformatics, *Bioinformatics* 25 (11) (2009) 1422–1423.
- [35] S. Cohen-Boulakia, O. Biton, S. Davidson, C. Froidevaux, *Bioguides*: querying multiple sources with a user-centric perspective, *Bioinformatics* 23 (10) (2007) 1301–1303.
- [36] T. Hulsen, S.S. Jamuar, A.R. Moody, J.H. Karnes, O. Varga, S. Hedensted, R. Spreafico, D.A. Hafler, E.F. McKinney, From big data to precision medicine, *Front. Med.* 6 (2019) 34.
- [37] T.R. Ross, D. Ng, J.S. Brown, R. Pardee, M.C. Hornbrook, G. Hart, J.F. Steiner, The hmo research network virtual data warehouse: a public data model to support collaboration, *eGEMs (Generating Evidence & Methods to improve patient outcomes)* 2 (1) (2014).
- [38] S.T. Savitz, L.A. Savitz, N.S. Fleming, N.D. Shah, A.S. Go, How much can we trust electronic health record data? in: *Healthcare*, Elsevier, 2020, p. 100444.
- [39] H. Spengler, I. Gatz, F. Kohlmayer, K.A. Kuhn, F. Prasser, Improving data quality in medical research: monitoring architecture for clinical and translational data warehouses, in: *2020 IEEE 33rd International Symposium on Computer-Based Medical Systems (CBMS)*, IEEE, 2020, pp. 415–420.
- [40] N.G. Weiskopf, C. Weng, Methods and dimensions of electronic health record data quality assessment: enabling reuse for clinical research, *J. Am. Med. Inform. Assoc.* 20 (1) (2013) 144–151.
- [41] G.O. Consortium, The gene ontology resource: 20 years and still going strong, *Nucl. Acids Res.* 47 (D1) (2019) D330–D338.
- [42] S. de Coronado, L.W. Wright, G. Fragoso, M.W. Haber, E.A. Hahn-Dantona, F.W. Hartel, S.L. Quan, T. Safran, N. Thomas, L. Whiteman, The nci thesaurus quality assurance life cycle, *J. Biomed. Inform.* 42 (3) (2009) 530–539.
- [43] R. Cruz Correia, D. Ferreira, G. Baceral, P. Marques, P. Maranhão, Personalised medicine challenges: quality of data, *Int. J. Data Sci. Anal.* 6 (3) (2018) 251–259.
- [44] M. Courtot, L. Cherubin, A. Faulconbridge, D. Vaughan, M. Green, D. Richardson, P. Harrison, P.L. Whetzel, H. Parkinson, T. Burdett, Biosamples database: an updated sample metadata hub, *Nucl. Acids Res.* 47 (D1) (2018) D1172–D1178.
- [45] Y. Cui, J. Widom, J.L. Wiener, Tracing the lineage of view data in a warehousing environment, *ACM Trans. Database Syst.* 25 (2) (2000) 179–227.
- [46] F. Cumbo, G. Fison, S. Ceri, M. Masseroli, E. Weitschek, TCGA2BED: extracting, extending, integrating, and querying the cancer genome atlas, *BMC Bioinform.* 18 (1) (2017) 6.
- [47] H.-H. Do, E. Rahm, Flexible integration of molecular-biological annotation data: the GenMapper approach, in: *International Conference on Extending Database Technology*, Springer, 2004, pp. 811–822.
- [48] The ENCODE Project Consortium, Overall coordination (data analysis coordination), Dunham, I. et al. An integrated encyclopedia of DNA elements in the human genome, *Nature* 489 (2012) 57–74, doi:10.1038/nature11247.
- [49] L. Etcheverry, A. Marotta, R. Ruggia, Data quality metrics for genome wide association studies, in: *2010 Workshops on Database and Expert Systems Applications*, IEEE, 2010, pp. 105–109.
- [50] W. Fan, Data quality: from theory to practice, *ACM SIGMOD Record* 44 (3) (2015) 7–18.
- [51] A. Frankish, M. Diekhans, A.-M. Ferreira, R. Johnson, I. Jungreis, J. Loveland, J.M. Mudge, C. Sis, J. Wright, J. Armstrong, et al., GENCODE reference annotation for the human and mouse genomes, *Nucl. Acids Res.* 47 (D1) (2018) D766–D773.
- [52] I. Gabdank, E.T. Chan, J.M. Davidson, J.A. Hilton, C.A. Davis, U.K. Baymuradov, A. Narayanan, K.C. Onate, K. Graham, S.R. Miyasato, et al., Prevention of data

- duplication for high throughput sequencing repositories, Database 2018 (2018) bay008.
- [53] E. Galeota, K. Kishore, M. Pelizzola, Ontology-driven integrative analysis of omics data through onassis, *Sci. Rep.* 10 (1) (2020) 1–9.
- [54] E. Galeota, M. Pelizzola, Ontology-based annotations and semantic relations in large-scale (epi)genomics data, *Brief. Bioinform.* 18 (3) (2017) 403–412.
- [55] E. Gasteiger, E. Jung, A.M. Bairoch, Swiss-prot: connecting biomolecular knowledge via a protein database, *Curr. Issues Mol. Biol.* 3 (3) (2001) 47–55.
- [56] M. Ghandi, F.W. Huang, J. Jané-Valbuena, G.V. Kryukov, C.C. Lo, E.R. McDonald, J. Barretina, E.T. Gelfand, C.M. Bielski, H. Li, et al., Next-generation characterization of the cancer cell line encyclopedia, *Nature* 569 (7757) (2019) 503.
- [57] C.B. Giles, C.A. Brown, M. Ripperger, Z. Dennis, X. Roonnarinesingh, H. Porter, A. Perz, J.D. Wren, Ale: automated label extraction from geo metadata, *BMC Bioinform.* 18 (14) (2017) 509.
- [58] C. Goble, R. Stevens, D. Hull, K. Wolstencroft, R. Lopez, Data curation+ process curation= data integration+ science, *Brief. Bioinform.* 9 (6) (2008) 506–517.
- [59] R.S. Gonçalves, M.A. Musen, The variable quality of metadata about biological samples used in biomedical experiments, *Sci. Data* 6 (2019) 190021.
- [60] A.J. Gray, C.A. Goble, R. Jimenez, Bioschemas: from potato salad to protein annotation, in: *International Semantic Web Conference (Posters, Demos & Industry Tracks)*, 2017.
- [61] J. Grosjean, T. Merabti, B. Dahamna, I. Kergourlay, B. Thirion, L.F. Soualmia, S.J. Darmoni, et al., Health multi-terminology portal: a semantic added-value for patient safety, *Stud. Health Technol. Inform.* 166 (66) (2011) 129–138.
- [62] R.L. Grossman, A.P. Heath, V. Ferretti, H.E. Varmus, D.R. Lowy, W.A. Kibbe, L.M. Staudt, Toward a shared vision for cancer genomic data, *N. Engl. J. Med.* 375 (12) (2016) 1109–1112.
- [63] É. Guerin, G. Marquet, A. Burgun, O. Loréal, L. Berti-Équille, U. Leser, F. Mousouni, Integrating and warehousing liver gene expression data and related biomedical resources in GEDAW, in: *International Workshop on Data Integration in the Life Sciences*, Springer, 2005, pp. 158–174.
- [64] C. Hedeler, P. Missier, Information quality management challenges for high-throughput data, *Biol. Database Model.* (2008) 81.
- [65] S. Henley-Smith, D. Boyle, K. Gray, Improving a secondary use health data warehouse: Proposing a multi-level data quality framework, eGEMs (Generating Evidence & Methods to improve patient outcomes) 7 (1) (2019).
- [66] C.-C. Huang, Z. Lu, Community challenges in biomedical text mining over 10 years: success, failure and the future, *Brief. Bioinform.* 17 (1) (2016) 132–144.
- [67] W. Huber, V.J. Carey, R. Gentleman, S. Anders, M. Carlson, B.S. Carvalho, H.C. Bravo, S. Davis, L. Gatto, T. Girke, et al., Orchestrating high-throughput genomic analysis with bioconductor, *Nat. Methods* 12 (2) (2015) 115–121.
- [68] J. Ison, H. Ienasescu, P. Chmura, E. Rydzka, H. Menager, M. Kalaš, V. Schwämmle, B. Grüning, N. Beard, R. Lopez, et al., The bio.tools registry of software tools and data resources for the life sciences, *Genome Biol.* 20 (1) (2019) 1–4.
- [69] J. Ison, M. Kalaš, I. Jonassen, D. Bolser, M. Uludag, H. McWilliam, J. Malone, R. Lopez, S. Pettifer, P. Rice, Edam: an ontology of bioinformatics operations, types of data and identifiers, topics and formats, *Bioinformatics* 29 (10) (2013) 1325–1332.
- [70] H. Ji, R.W. Davis, Data quality in genomics and microarrays, *Nat. Biotechnol.* 24 (9) (2006) 1112–1113.
- [71] C. Jonquet, N.H. Shah, M.A. Musen, The open biomedical annotator, *Summit Transl. Bioinform.* 2009 (2009) 56.
- [72] M.G. Kahn, T.J. Callahan, J. Barnard, A.E. Bauck, J. Brown, B.N. Davidson, H. Estiri, C. Goerg, E. Holve, S.G. Johnson, et al., A harmonized data quality assessment terminology and framework for the secondary use of electronic health record data, eGEMs (Generating Evidence & Methods to improve patient outcomes) 4 (1) (2016).
- [73] Y. Kodama, M. Shumway, R. Leinonen, The sequence read archive: explosive growth of sequencing data, *Nucl. Acids Res.* 40 (D1) (2011) D54–D56.
- [74] J.L. Koh, M.L. Lee, A.M. Khan, P.T. Tan, V. Brusic, Duplicate detection in biological data using association rule mining, in: *Proceedings of the second European Workshop on Data Mining and text Mining in Bioinformatics*, Citeseer, 2004, pp. 35–41.
- [75] A. Kundaje, W. Meuleman, J. Ernst, M. Bilenky, A. Yen, A. Heravi-Moussavi, P. Kheradpour, Z. Zhang, J. Wang, M.J. Ziller, et al., Integrative analysis of 111 reference human epigenomes, *Nature* 518 (7539) (2015) 317.
- [76] A. León, J. Reyes, V. Burriel, F. Valverde, Data quality problems when integrating genomic information, in: *International Conference on Conceptual Modeling*, Springer, 2016, pp. 173–182.
- [77] J. Lonsdale, J. Thomas, M. Salvatore, R. Phillips, E. Lo, S. Shad, R. Hasz, G. Walters, F. Garcia, N. Young, et al., The genotype-tissue expression (GTEx) project, *Nat. Genet.* 45 (6) (2013) 580.
- [78] D. Maglott, J. Ostell, K.D. Pruitt, T. Tatusova, Entrez gene: gene-centered information at NCBI, *Nucl. Acids Res.* 39 (suppl_1) (2010) D52–D57.
- [79] J. Malone, E. Holloway, T. Adamusiak, M. Kapushesky, J. Zheng, N. Kolesnikov, A. Zhukova, A. Brazma, H. Parkinson, Modeling sample variables with an experimental factor ontology, *Bioinformatics* 26 (8) (2010) 1112–1118.
- [80] A. Martinez, J. Hammer, Making quality count in biological data sources, in: *Proceedings of the 2nd international workshop on Information quality in information systems*, 2005, pp. 16–27.
- [81] A. Martinez, J. Hammer, S. Ranka, Biodq: data quality estimation and management for genomics databases, in: *International Symposium on Bioinformatics Research and Applications*, Springer, 2008, pp. 469–480.
- [82] M. Martínez-Romero, M.J. O'Connor, A.L. Egyedi, D. Willrett, J. Hardi, J. Graybeal, M.A. Musen, Using association rule mining and ontologies to generate metadata recommendations from multiple biomedical databases, *Database* (2019).
- [83] M. Masseroli, A. Canakoglu, P. Pinoli, A. Kaitoua, A. Gulino, O. Horlova, L. Nanni, A. Bernasconi, S. Perna, E. Stamoulakatou, S. Ceri, Processing of big heterogeneous genomic datasets for tertiary analysis of next generation sequencing data, *Bioinformatics* 35 (5) (2018) 729–736.
- [84] M. Masseroli, A. Kaitoua, P. Pinoli, S. Ceri, Modeling and interoperability of heterogeneous genomic big data for integrative processing and querying, *Methods* 111 (2016) 3–11.
- [85] J.A. McMurry, N. Juty, N. Blomberg, T. Burdett, T. Conlin, N. Conte, M. Courtot, J. Deck, M. Dumontier, D.K. Fellows, et al., Identifiers for the 21st century: How to design, provision, and reuse persistent identifiers to maximize utility and impact of life science data, *PLoS Biol.* 15 (6) (2017).
- [86] T.F. Meehan, A.M. Masci, A. Abdulla, L.G. Cowell, J.A. Blake, C.J. Mungall, A.D. Diehl, Logical development of the cell ontology, *BMC Bioinform.* 12 (1) (2011) 6.
- [87] P. Missier, S. Embury, M. Greenwood, A. Preece, B. Jin, Quality views: capturing and exploiting the user perspective on data quality, in: *Proceedings of the International Conference on Very Large Data Bases*, 2006, pp. 977–988.
- [88] F. Mousouni, L. Berti-Équille, F. le Développement, in: M. Elloumi, A.Y. Zomaya (Eds.), *Cleaning, integrating, and warehousing genomic data from biomedical resources*, Hoboken, NJ: John Wiley and Sons, Inc., 2013, pp. 35–58.
- [89] F. Mousouni, L. Berti-Equille, G. Rozé, O. Loréal, E. Guérin, Qdex: a database profiler for generic bio-data exploration and quality aware integration, in: *International Conference on Web Information Systems Engineering*, Springer, 2007, pp. 5–16.
- [90] H. Müller, J.-C. Freytag, U. Leser, Improving data quality by source analysis, *J. Data Inf. Qual.* 2 (4) (2012) 1–38.
- [91] H. Müller, F. Naumann, Data quality in genome databases, in: *Eighth International Conference on Information Quality (ICIQ 2003)*, November 7-9, 2003, 2003, pp. 269–284.
- [92] C.J. Mungall, C. Toornia, G.V. Gkoutos, S.E. Lewis, M.A. Haendel, Uberon, an integrative multi-species anatomy ontology, *Genome Biol.* 13 (1) (2012) R5.
- [93] M.A. Musen, S.-A. Sansone, K.-H. Cheung, S.H. Kleinstein, M. Crafts, S.C. Schürer, J. Graybeal, Cedar: Semantic web technology to support open science, in: *Companion Proceedings of the The Web Conference 2018*, International World Wide Web Conferences Steering Committee, 2018, pp. 427–428.
- [94] S. Neph, M.S. Kuehn, A.P. Reynolds, E. Haugen, R.E. Thurman, A.K. Johnson, E. Rynes, M.T. Maurano, J. Vierstra, S. Thomas, et al., Bedops: high-performance genomic feature operations, *Bioinformatics* 28 (14) (2012) 1919–1920.
- [95] N.A. O’Leary, M.W. Wright, J.R. Brister, S. Ciuffo, D. Haddad, R. McVeigh, B. Rajput, B. Robbertse, B. Smith-White, D. Ako-Adjei, et al., Reference sequence (RefSeq) database at NCBI: current status, taxonomic expansion, and functional annotation, *Nucl. Acids Res.* 44 (D1) (2015) D733–D745.
- [96] D. Oliveira, A.S. Butt, A. Haller, D. Rebolz-Schuhmann, R. Sahay, Where to search top-k biomedical ontologies? *Brief. Bioinform.* 20 (4) (2019) 1477–1491.
- [97] A.L. Palacio, Ó.P. López, J.C.C. Ródenas, A method to identify relevant genome data: conceptual modeling for the medicine of precision, in: *International Conference on Conceptual Modeling*, Springer, 2018, pp. 597–609.
- [98] Y.M. Park, S. Squizzato, N. Buso, T. Gur, R. Lopez, The EBI search engine: EBI search as a service-making biological data accessible for all, *Nucl. Acids Res.* 45 (W1) (2017) W545–W549.
- [99] O. Pastor, A.P. León, J.F.R. Reyes, A.S. García, J.C.R. Casamayor, Using conceptual modeling to improve genome data management, *Brief. Bioinform.* (2020), doi:10.1093/bib/bbaa100.
- [100] L. Posch, M. Panahiazar, M. Dumontier, O. Gevaert, Predicting structured metadata from unstructured metadata, *Database* (2016).
- [101] S. Poux, C.N. Arighi, M. Magrane, A. Bateman, C.-H. Wei, Z. Lu, E. Boutet, H. Bye-A-Jee, M.L. Famiglietti, B. Roechert, et al., On expert curation and scalability: Uniprotkb/swiss-prot as a case study, *Bioinformatics* 33 (21) (2017) 3454–3460.
- [102] S. Poux, M. Magrane, C.N. Arighi, A. Bridge, C. O’Donovan, K. Laiho, U. Consortium, et al., Expert curation in uniprotkb: a case study on dealing with conflicting and erroneous data, *Database* (2014).
- [103] A.R. Quinlan, I.M. Hall, Bedtools: a flexible suite of utilities for comparing genomic features, *Bioinformatics* 26 (6) (2010) 841–842.
- [104] N.S. Rajan, R. Gouripeddi, P. Mo, R.K. Madsen, J.C. Facelli, Towards a content agnostic computable knowledge repository for data quality assessment, *Comput. Methods Programs Biomed.* 177 (2019) 193–201.
- [105] T.C. Redman, A. Blanton, *Data Quality for the Information Age*, Artech House, Inc., 1996.
- [106] J.F.R. Román, Ó. Pastor, J.C. Casamayor, F. Valverde, Applying conceptual modeling to better understand the human genome, in: *International Conference on Conceptual Modeling*, Springer, 2016, pp. 404–412.
- [107] S.A. Samarajiva, I. Olan, D. Bihary, Challenges and cases of genomic data integration across technologies and biological scales, in: *Advanced Data Analytics in Health*, Springer, 2018, pp. 201–216.
- [108] K. Sanderson, Bioinformatics: curation generation, *Nature* 470 (7333) (2011) 295–296.
- [109] S.-A. Sansone, P. McQuilton, P. Rocca-Serra, A. Gonzalez-Beltran, M. Izzo, A.L. Lister, M. Thurston, Fairsharing as a community approach to standards, repositories and policies, *Nat. Biotechnol.* 37 (4) (2019) 358–367.
- [110] E.W. Sayers, M. Cavanaugh, K. Clark, J. Ostell, K.D. Pruitt, I. Karsch-Mizrachi, GenBank, *Nucl. Acids Res.* 47 (D1) (2019) D94–D99.
- [111] S.C. Schuster, Next-generation sequencing transforms today’s biology, *Nat. Methods* 5 (1) (2008) 16–18.
- [112] N.H. Shah, C. Jonquet, A.P. Chiang, A.J. Butte, R. Chen, M.A. Musen, Ontology-driven indexing of public datasets for translational bioinformatics, *BMC Bioinform.* (10 Suppl 2) (2009) S1.

- [113] D. Soh, D. Dong, Y. Guo, L. Wong, Consistency, comprehensiveness, and compatibility of pathway databases, *BMC Bioinform.* 11 (1) (2010) 449.
- [114] Z.D. Stephens, S.Y. Lee, F. Faghri, R.H. Campbell, C. Zhai, M.J. Efron, R. Iyer, M.C. Schatz, S. Sinha, G.E. Robinson, Big data: astronomical or genomics? *PLoS Biol.* 13 (7) (2015) e1002195.
- [115] B. Stvilia, L. Gasser, M.B. Twidale, L.C. Smith, A framework for information quality assessment, *J. Am. Soc. Inf. Sci. Technol.* 58 (12) (2007) 1720–1733.
- [116] J.G. Tate, S. Bamford, H.C. Jubb, Z. Sondka, D.M. Beare, N. Bindal, H. Boutselakis, C.G. Cole, C. Creatore, E. Dawson, et al., COSMIC: the catalogue of somatic mutations in cancer, *Nucl. Acids Res.* 47 (D1) (2018) D941–D947.
- [117] N. Vasilevsky, T. Johnson, K. Corday, C. Torniai, M. Brush, E. Segerdell, M. Wilson, C. Shaffer, D. Robinson, M. Haendel, Research resources: curating the new eagle-i discovery system, *Database* (2012).
- [118] A.K. Veiga, A.M. Saraiva, A.D. Chapman, P.J. Morris, C. Gendreau, D. Schigel, T.J. Robertson, A conceptual framework for quality assessment and management of biodiversity data, *PLoS One* 12 (6) (2017).
- [119] R.Y. Wang, D.M. Strong, Beyond accuracy: what data quality means to data consumers, *J. Manag. Inf. Syst.* 12 (4) (1996) 5–33.
- [120] Z. Wang, A. Lachmann, A. Ma'ayan, Mining data and metadata from the gene expression omnibus, *Biophys. Rev.* 11 (1) (2019) 103–110.
- [121] J.N. Weinstein, E.A. Collisson, G.B. Mills, K.R.M. Shaw, B.A. Ozenberger, K. Ellrott, I. Shmulevich, C. Sander, J.M. Stuart, C.G.A.R. Network, et al., The cancer genome atlas pan-cancer analysis project, *Nat. Genet.* 45 (10) (2013) 1113.
- [122] M.D. Wilkinson, M. Dumontier, I.J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L.B. da Silva Santos, P.E. Bourne, et al., The fair guiding principles for scientific data management and stewardship, *Sci. Data* 3 (2016).
- [123] B. Yates, B. Braschi, K.A. Gray, R.L. Seal, S. Tweedie, E.A. Bruford, Genenames.org: the hgnc and vnc resources in 2017, *Nucl. Acids Res.* (2016) gkw1033.
- [124] A. Zaveri, W. Hu, M. Dumontier, Metacrowd: crowdsourcing biomedical metadata quality assessment, *Hum. Comput.* 6 (1) (2019) 98–112.
- [125] D.R. Zerbino, P. Achuthan, W. Akanni, M.R. Amode, D. Barrell, J. Bhai, K. Billis, C. Cummins, A. Gall, C.G. Girón, et al., Ensembl 2018, *Nucl. Acids Res.* 46 (D1) (2017) D754–D761.
- [126] J. Zhang, J. Baran, A. Cros, J.M. Guberman, S. Haider, J. Hsu, Y. Liang, E. Rivkin, J. Wang, B. Whitty, et al., International cancer genome consortium data portal—a one-stop shop for cancer genomics data, *Database* (2011).
- [127] R. Zheng, C. Wan, S. Mei, Q. Qin, Q. Wu, H. Sun, C.-H. Chen, M. Brown, X. Zhang, C.A. Meyer, et al., Cistrome Data Browser: expanded datasets and new tools for gene regulatory analysis, *Nucl. Acids Res.* 47 (D1) (2018) D729–D735.



Anna Bernasconi is a post-doctoral researcher at Politecnico di Milano. Her interests include bioinformatics data and metadata integration methodologies to support biological queries. She focuses on open data repositories and biomedical ontologies.