

META-BASE: a Novel Architecture for Large-Scale Genomic Metadata Integration

Anna Bernasconi, Arif Canakoglu, Marco Masseroli and Stefano Ceri

Abstract—The integration of genomic metadata is, at the same time, an important, difficult, and well-recognized challenge. It is important because a wealth of public data repositories is available to drive biological and clinical research; combining information from various heterogeneous and widely dispersed sources is paramount to a number of biological discoveries. It is difficult because the domain is complex and there is no agreement among the various metadata definitions, which refer to different vocabularies and ontologies. It is well-recognized in the bioinformatics community because, in the common practice, repositories are accessed one-by-one, learning their specific metadata definitions as result of long and tedious efforts, and such practice is error-prone. In this paper, we describe META-BASE, an architecture for integrating metadata extracted from a variety of genomic data sources, based upon a structured transformation process. We present a variety of innovative techniques for data extraction, cleaning, normalization and enrichment. We propose a general, open and extensible pipeline that can easily incorporate any number of new data sources, and propose the resulting repository – already integrating several important sources – which is exposed by means of practical user interfaces to respond biological researchers' needs.

Index Terms—Data Integration, Genomic Datasets, Metadata Management, Open Data, Rule-Based Languages, Bioinformatics.

1 INTRODUCTION

GENOMIC research is showing a variety of initiatives for the production of high-value biological and clinical datasets, stored in open repositories and available to the research community for secondary research use. Some examples include the Encyclopedia of DNA Elements (ENCODE [1]), The Cancer Genome Atlas (TCGA [2]) and its successor Genomic Data Commons (GDC [3]), Roadmap Epigenomics Project [4], 1000 Genomes [5], and many others. Metadata, generally defined as “data about data”, in genomics repositories collectively refers to the descriptions of experimental and annotation datasets; it specifies the experimental conditions, the cell lines or tissues, the donors with their demography, phenotypes, and treatments, and the process of extraction of stored genomic signals with the used technological devices. By inspecting metadata, it is possible to locate the datasets that better fit for formulating queries over the genome; these in turn can answer important questions in modern biology and precision medicine.

Unfortunately, while we observe a good convergence in the definition of data formats and protocols for genomic information (e.g., BED, MAF, VCF, GTF), no agreement for a common metadata format has been reached so far: metadata of distinct repositories often disagree on their entities, attributes and values, and have no associated conceptual representations [6], [7]. In earlier work, we developed a conceptual approach to metadata integration and presented the Genomic Conceptual Model (GCM), which mediates over the most important and complex data sources [8]. This paper is focused on the process required to generate the GCM content and on the resulting repository. First, we describe **META-BASE**, a novel architecture for the integration

of genomic datasets; it is deployed as a generic pipeline of six progressive steps for data integration, applicable to arbitrary genomic data sources providing semi-structured metadata descriptions. Four steps are driven by source-specific modules, the others are source-independent. Two steps are assisted by tools that help the designer in the progressive creation and adaptation of data management rules, with the general objective of minimizing the cognitive effort required from integration designers. Then, we describe the pipeline that generates the **META-BASE repository**, a very large integrated repository of tertiary genomic datasets. In this paper, we focus on the integration of three data sources featuring complex metadata: ENCODE, GDC, and Roadmap Epigenomics. In addition, META-BASE integrates other sources, whose conceptual complexity is much simpler: genomic annotation data from GENCODE [9] and RefSeq [10]; variation data from 1000 Genomes; topologically associating domains (TADs) [11] from GEO repository [12]; epigenomic experiments metadata from Cistrome [13]. The META-BASE repository will continue to grow in the next years, responding to biological and clinical needs.

Every step of the META-BASE pipeline produces a data ingestion program that can be applied to data sources after an initial design; these programs need to be adapted only in case of structural changes of the data sources. The process is extensible, as the designer who wants to add a new source has just to add new definitions and rules to the data integration framework.

Within the data enrichment step of the META-BASE pipeline, we also use selected ontological sources for improving value matching, which is extended from exact match to semantic match inclusive of the use of synonyms, hyponyms and hyperonyms; they enable simple value conversion strategies, which capture some value mismatches that may occur in different repositories.

- *The authors are with Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Via Ponzio 34/5, 20133, Italy. E-mail: first.last@polimi.it*

Queries upon the META-BASE repository are supported by GenoSurf [14], a user-friendly, attribute-based and keyword-based search interface; the metadata attributes supported by GenoSurf are organized using the GCM layout [15]. Queries upon the META-BASE repository can be used for producing as a result the Uniform Resource Identifiers (URIs) of the relevant data in the source repositories; scientists can build over them an arbitrary genomic computation, using any bioinformatics system and resource. In this way, the META-BASE repository provides a conceptual entry point to the supported genomic data sources. In addition, the META-BASE pipeline and repository feed an architecture for genomic data processing, defined in [16], providing portable and scalable genomic data management on powerful servers and clusters¹; in such distinct environment, metadata can be queried together with their respective datasets using GenoMetric Query Language (GMQL), a high-level domain-specific query language [17].

The most innovative aspects of our work are: from a computer science perspective, the design and engineering of an end-to-end pipeline whose steps make novel use of rewrite rules for data cleaning, mapping, normalization, enrichment and integrity verification; from a biological perspective, the partitioning schemes for each data source and the selection of the ontologies providing enrichment for specific GCM attributes; from an application perspective, the provisioning to the broad community of computational biology and bioinformatics of new interfaces and systems fed by a solid data architecture.

Paper organization. Section 2 overviews the GCM, which is at the base of this work, and provides a motivating example. Section 3 describes our generalized approach to metadata management and integration. Section 4 describes the pipeline to extract metadata from original sources and to prepare it for integration. Section 5 shows the integration process towards the final META-BASE repository, which includes ontological enrichment. Section 6 discusses the effectiveness of our approach. Section 7 describes the architecture of the system. Section 8 overviews related work, and Section 9 mentions future developments and concludes.

2 BACKGROUND

In [8] we originally presented the Genomic Conceptual Model (GCM), an Entity-Relationship model used to describe metadata of genomic data sources; its current version² is shown in Fig. 1. The main objective of GCM is to recognize the common organization for a limited set of concepts that are supported by most genomic data sources, although with very different names and formats. GCM is organized as a star-schema [18] – typically used in the design of data warehouses [19], consisting of a *fact* table referencing any number of *dimension* tables – centered on the ITEM entity (i.e., the *fact*, in the data warehouses' jargon), representing an elementary experimental file of genomic regions and their attributes. Files are typically used by biologists for data extraction, analysis and visualization operations. Four

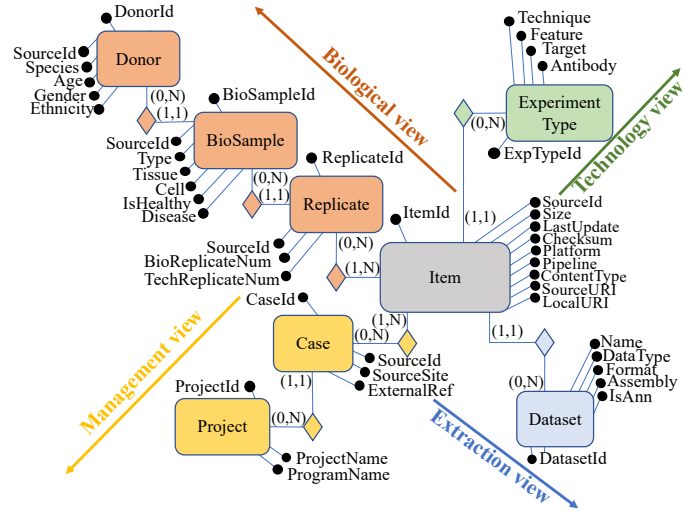


Fig. 1. Genomic Conceptual Model (GCM).

hierarchical *dimensions* describe: 1) the *biological* elements involved in the experiment: the sequenced sample, its preparation, its donor; 2) the *technology* used in the experiment, including the specific technique; 3) the *management* aspects of the experiment: the projects/organizations behind its preparation and production; 4) the *extraction* parameters used for internal selection and organization of items.

The **Central Entity** ITEM features the *SourceId* attribute, which identifies it uniquely on the source system, along with its *Size*, *LastUpdate*, and *Checksum*. *Platform* and *Pipeline* are respectively used to provide references to the methods and parameters used for production and processing of sequenced raw data and its processed data. *ContentType* accepts values such as “peaks”/“hotspots”/“exon quantifications” when the contained regions are experimental, or “gene”/“transcript”/“promoter” when they are annotations. The ITEM is physically available for download at the *LocalURI* and, in its original form, at the *SourceURI*.

The **Biological View** consists of the chain of entities ITEM-REPLICATE-BIOSAMPLE-DONOR, representing the biological elements that contribute to the ITEM production. An ITEM is associated with one or more REPLICATES, each originated by a BIOSAMPLE, each derived from a DONOR. DONOR is identified by a *SourceId*. It represents an individual (characterized by *Age*, *Gender* and *Ethnicity*) or strain of a specific organism (*Species*) from which the biological material was derived or the cell line was established. BIOSAMPLE is identified within possibly multiple original sources by *SourceId*. Its characterizing *Type* expresses values such as “cell line”/“tissue”/“primary cell”, depending on the kind of material sample used for the experiment. *Cell* includes information of (single) cells in their natural state, immortalized cell lines, or cells differentiated from specific cell types. *Tissue* regards a multicellular component in its natural state, or the provenance tissue of the *Cell(s)* of a biosample. *IsHealthy* denotes a healthy (normal/control) or non-healthy (e.g., tumoral) sample, and *Disease* stores information about the disease investigated with the sample. REPLICATE is useful to model cases where an assay is performed multiple times on similar biological material. If repeated on separate biological samples, the generated items are biological replica of a same experiment; if repeated

1. Based on Apache Spark <http://spark.apache.org/>
 2. Several adaptations of the model w.r.t. what was presented in [8] have been performed afterwards.

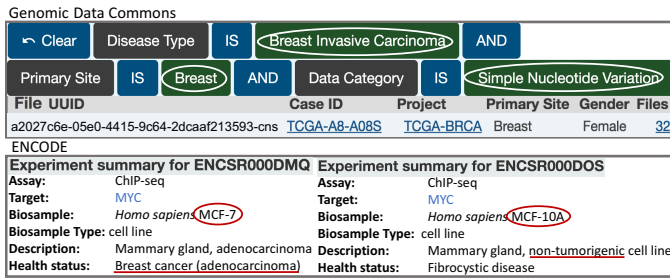


Fig. 2. Example of Web interfaces of data sources: GDC and ENCODE.

on two portions from the same biological sample (treated for example with same growth, excision, and knockdown), the items are technical replicates. This occurs only in some epigenomic data sources (such as ENCODE and Roadmap Epigenomics) that perform assay replication.

The **Technology View** describes the technology used to produce the data ITEM. An ITEM is associated by means of a one-to-many relationship with a given EXPERIMENTTYPE, which includes the *Technique* (e.g., “ChIP-seq”, “DNase-seq”, “RRBS”) and the *Feature*, which denotes the specific genomic aspect described by the experiment (e.g., “Copy Number Variation”, “Histone Modification”, “Transcription Factor”). When the *Technique* is “ChIP-seq”, *Target* and *Antibody* are needed to further characterize the experiment.

The **Management View** consists of the chain of entities ITEM-CASE-PROJECT describing the organizational process for the production of items. CASE represents a set of items that have been collected within the same research study. *SourceId* and *ExternalRef* contain identifiers respectively taken from the main original source and alternative sources. *SourceSite* represents the physical site where the material is analyzed. PROJECT represents a project, a program, or a single initiative responsible for the production of the ITEM. It provides a single point of reference to find diverse data types generated in a same research context.

The **Extraction View** includes the entity DATASET, used to describe common properties of homogeneous items. Its attributes include a *Name*, useful to locate and organize data, the *DataType*, describing the specific kind of genomic data contained in the items of such dataset (e.g., “peaks”, “copy number segments”, “gene expression quantification”), the *Format*, which denotes the ITEM data file format (e.g., “bed”, or more specific ones such as “narrowPeak” and “broadPeak”), and the reference genome alignment (*Assembly*). *IsAnn* distinguishes between experimental items (describing arbitrary genomic regions) and annotations (describing known genomic regions).

2.1 Motivating Example

To motivate our effort, we introduce an example that simulates the research of data suitable for a genomics project on two different sources. We focus on a simple situation, which can be appreciated even by a reader with limited biological background. Consider a comparison study between a human non-healthy breast tissue, affected by carcinoma, and a healthy sample coming from the same tissue type. A researcher in the field locates two portals having interest-

ing data for this analysis. The results obtained after some browsing are reported in Fig. 2.

For the healthy data, the chosen source is GDC Data Portal, an important repository of human cancer genomic data. As it can be seen on the top of Fig. 2, typically more data files can be retrieved by composing a query that allows locating variation data on “Breast Invasive Carcinoma” from “Breast” tissue. By browsing several metadata information sections (sometimes hard to identify), the researcher can find files corresponding to “normal” (i.e., non-tumoral) tissue.

To compare such data with others from a diseased reference, the researcher considers additional datasets coming from cell lines, i.e., cell cultures that have been permanently established and made immortal. Cell lines are frequently used in place of primary cells to study biological processes, as the scientific community tends to accept the derived findings more readily. On ENCODE, the researcher chooses both a tumor cell line (bottom left of Fig. 2) and a normal cell line (bottom right of Fig. 2), to make a control comparison. “MCF-7” is a cell line from a diseased tissue affected by “Breast cancer (adenocarcinoma)”, while “MCF-10A” is its widely considered non-tumorigenic counterpart.

Note that some external knowledge is necessary in order to find these connections, which cannot be obtained on the mentioned portals. Regarding the *disease*, note that “Breast Invasive Carcinoma” and “breast cancer (adenocarcinoma)” are related sub-types of “breast carcinoma” (as observed in the EFO and DOID ontologies [20]); this allows us to compare GDC data with the dataset from ENCODE. For what concerns the *cell line*, researchers typically query specific databases (such as the COSMIC Cell Lines Browser³) or dedicated forums to discover tumor/normal matched cell line pairs. This information is not encoded in a unique way over data sources and is often missing.

3 APPROACH

The six phases of the META-BASE approach can be seen in Fig. 3. Through downloading, metadata is imported at the repository site. During transformation, they are translated into raw attribute-value pairs, which are then cleaned, thereby producing a collection of clean metadata pairs for each source. The mapper extracts information from these pairs and adds it to GCM; the values are then normalized (resorting to generic term IDs that may take specific sets of values) and enriched (by means of external ontologies). Finally, the consistency of the database content is checked with respect to integrity constraints.

For exemplifying the META-BASE framework, we consider three important and complex data sources:

- ENCODE contains datasets related to the functional DNA sequences, intervening at the protein/RNA levels, and to the regulatory elements that control gene expression.
- GDC contains datasets from TCGA program, related to many aspects of cancer genomics.
- Roadmap Epigenomics (indicated as ‘REP’ in the following) contains datasets related to epigenomic features in human normal tissues often involved in human diseases.

3. https://cancer.sanger.ac.uk/cell_lines/

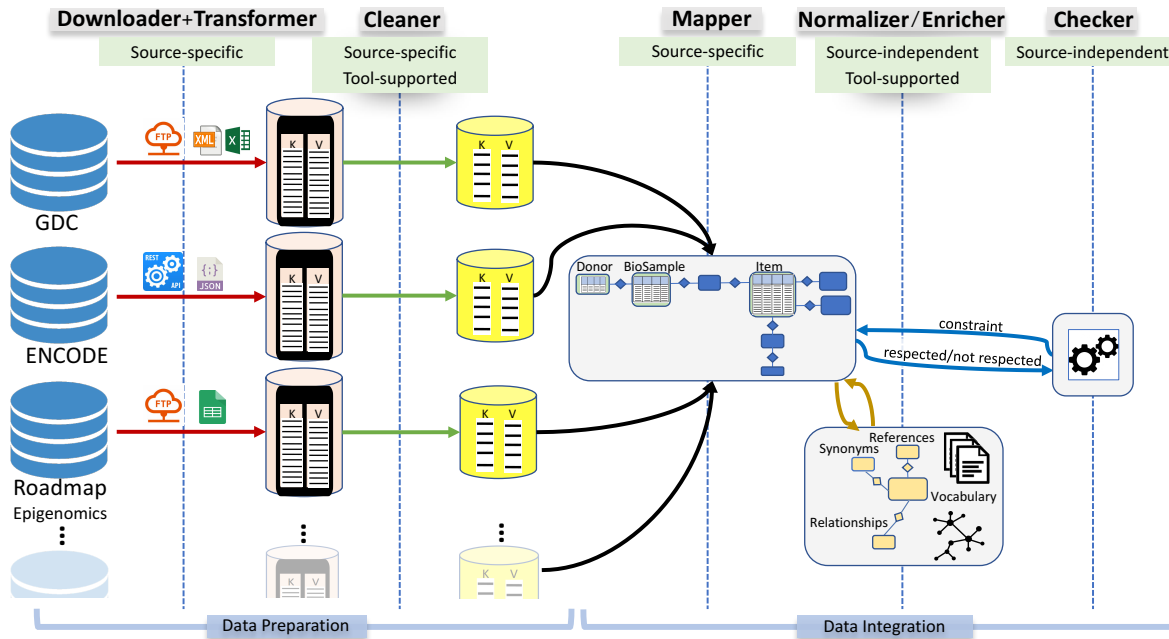


Fig. 3. The overall data preparation and integration process.

The above repositories are subject to rapid changes, as each source is a continuously evolving system. Luckily, most changes are additive and use already existing metadata in their descriptions. For this reason, we approach each source with an initial activity for the production of a source-specific set of metadata transformation rules, followed by periodic data integration sessions, where new items are discovered and their metadata are modeled.

At the same time, rules capture only a portion of the data integration semantics, as we allow for exceptions. Attributes that are not in common to most sources, while specific for few experiment types, are modeled as attribute-value pairs; related data is directly referenced from the ITEM entity.

4 DATA PREPARATION

The META-BASE data preparation pipeline allows us to extract metadata from a set of selected data sources and arrange it for integration. Metadata is first downloaded in their original formats (Section 4.1); then, it is transformed into a ⟨key, value⟩ equivalent form (Section 4.2); finally, since it is still *raw metadata*, it is exposed to a cleaning process that aims to improve raw attribute names and to filter irrelevant metadata (Section 4.3) before data integration.

4.1 Data Download

The *Downloader* module produces files both for the genomic data and its metadata, in original source-specific format, at the processing site hosting our repository; it must be programmed or adapted for each source. In most target sources, several protocols or APIs are made available for data downloading, but they do not share any standard for the metadata description or format (e.g., XML, JSON, tab-delimited). Some sources provide a metadata file for each experiment. In other cases, a single metadata file describes a collection of experiments; also in such a case, we produce multiple metadata files, one for each experimental data file.

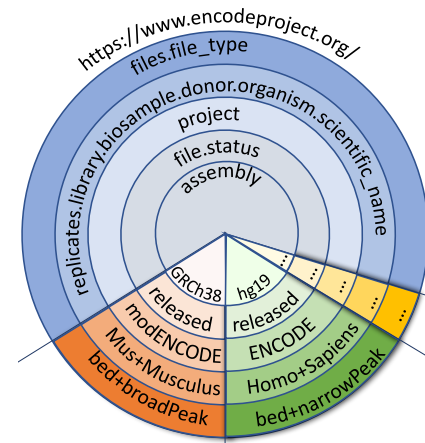


Fig. 4. Selection of portions from *ENCODE*. In the upper area we specify parameters names, in the two small bottom slices we specify example values, defining a partition of the source.

The main difficulty is to identify a specific data partitioning scheme at each source; in this way, each partition can be repeatedly accessed and source files that are added to or modified within the partition can be selectively recognized, avoiding the download for those source files that are unchanged. Table 1 illustrates three endpoints for data download used for the considered sources, specifying their protocol, request format, and example parameters for invocation. A partitioning scheme for the *ENCODE* data source is illustrated in Fig. 4, with a specific set of parameters used during download, corresponding to a partition.

Formalization. For a given source i , a *Downloader* is a method D_i for importing genomic data and metadata from a specific partition P_i of i . At each invocation of the method, a new set D_i of files (one for each data ITEM) is retrieved at the repository site, and associated with a signature ⟨*dataset_name*, *source*, *endpoint*, *parameters*⟩; parameters include the timestamp t_h of the download oper-

TABLE 1
 Endpoints for data download from sources and example invocations.

| | | |
|---------------|--|--|
| <i>ENCODE</i> | List of <code>file_id</code> Example <code>params</code> Download file | Protocol HTTP GET: <code>https://www.encodeproject.org/metadata/?type=Experiment&<params>/metadata.tsv</code> <code>assembly=hg19 & file.status=released & project=ENCODE & ... & files.file_type=bed+narrowPeak</code> <code>https://www.encodeproject.org/files/<file_id>/@@download/<file_id>.bed.gz</code> |
| <i>GDC</i> | List of <code>file_id</code> Example <code>params</code> Download file | Protocol HTTP POST: <code>https://api.gdc.cancer.gov/files</code> with <code><params></code> in Payload <code>field:cases.project.project_id-value:["TCGA-ACC"]</code> , <code>field:files.data_type-value:["Copy Number Segment"]</code> , ... <code>https://gdc-api.nci.nih.gov/data/<file_id></code> |
| <i>REP</i> | <code>dir</code> paths Example <code>dir</code> Download file | Protocol FTP: <code>http://egg2.wustl.edu/roadmap/data/byFileType/peaks/consolidated/<dir></code> <code>broadPeak</code> <code>http://egg2.wustl.edu/roadmap/data/byFileType/peaks/consolidated/<dir>/<file_name>.<dir>.gz</code> |

ation. In this way, future invocations of \mathcal{D}_i at time $t_k > t_h$ will be used to download information from P_i and then start a data integration session by tracking the changes that occurred to P_i at the data source between time t_h and t_k .

Method. Each download module first connects to the data source servers to retrieve the list of the identifiers of files that belong to the partition to download (corresponding to the GCM ITEMS). Many sources provide (semi)programmatic methods to translate a query composed on their portal visual interface into an API request or a downloadable list of files corresponding to the search; otherwise, this step has to be programmed.

For each ITEM, a downloader typically retrieves the *Size*, *LastUpdate* and *Checksum*, denoting properties of the data file; these are provided by most sources.⁴ We match these values with data that is stored in GCM, using the *SourceId* unique identifier. The matching allows us to pinpoint:

- New items: they are stored as genomic data files and their metadata are processed by invoking the pipeline discussed in this section.
- Matching items, having same *Size*, *LastUpdate* and *Checksum* values as their local values stored in the conceptual model: we reprocess just the metadata by invoking the pipeline discussed in this section, but we avoid the download of region data, which is typically much bigger in size. If any of the metadata values is different, we then download also the genomic data files.
- Missing items, i.e., items whose identifier was present at the previous invocation but it is no longer present: they are deprecated, the genomic data and metadata is copied to an archive that can only be inspected by archive lookups (but no longer retrieved by standard queries).

A downloader task splits originally downloaded metadata into files that correspond each to a single experimental region data file (e.g., ITEM). Eventually, we collect into the set of files \mathcal{D}_i all the metadata relative to new or changed items; these downloaded files are then used in the next steps of the pipeline. In parallel, the corresponding genomic data files are stored in the GMQL data repository (see Section 7).

Example. The *ENCODE* Web portal, described in [21], supports a faceted searching system that can be used to evaluate alternative options for metadata retrieval. Each search option produces different JSON objects. After careful analysis, we selected the option of retrieving the JSON file associated with an *experimental study* in “embedded” mode, as it includes compact information about all data files, replicates, and biosamples involved in it. Therefore, this was

4. If some of them are unavailable, we either compute them at the source or accept a less precise matching by using fewer parameters.

selected as the *ENCODE* metadata reference endpoint (see Table 1): when a downloader is invoked, it retrieves the data and metadata files partitioned by experiment.

As it can be observed in Listing 1, the information associated with the specific experiment with accession ENCSR635OSG is a hierarchically structured JSON file, including several embedded elements: information about the whole experimental study, arrays of “files” elements (a list of items included in the experimental study) and of “replicates” elements, along with other information.

After retrieval, the identifiers of the items belonging to the considered partition are recorded together with their size, last update date and checksum. Then, a downloader task separates the information retrieved for an experiment into several metadata files, each containing the information about a specific item of the experimental study. From the excerpt of Listing 1, two JSON files are created for items ENCF134AVY and ENCF429VMY, where the former one has one replica while the latter one has two. In this way, all following data preparation steps apply to metadata files that are in one-to-one correspondence with data ITEMS.

```
{
  "accession": "ENCSR635OSG",
  "assembly": ["hg19"],
  "award": {
    "pi": {
      "lab": {"name": "michael-snyder", ...},
      ...
    },
    ...
  },
  "dbxrefs": [],
  "files": [
    {
      "accession": "ENCF134AVY",
      "biological_replicates": [1], ...
    },
    {
      "accession": "ENCF429VMY",
      "biological_replicates": [1, 2],
      "file_type": "bed narrowPeak", ...
    },
    ...
  ],
  "replicates": [
    {
      "@id": "/replicates/4874c170-7124-4822-a058-4bb/",
      "biological_replicate_number": 1,
      "library": {
        "biosample": {
          "donor": {"age": "6", ...},
          "health_status": "healthy",
          ...
        },
        ...
      },
      "antibody": {"lot_id": "940739", ...},
      ...
    },
    {
      "@id": "/replicates/d42ff80d-67fd-45ee-9159-25a/",
      "biological_replicate_number": 2,
      "library": {
        "biosample": {
          "donor": {"age": "32", ...},
          "health_status": "healthy with non-
            obstructive coronary artery disease",
          ...
        },
        ...
      },
      "antibody": {"lot_id": "940739", ...},
      ...
    }
  ],
  ...
}
```

Listing 1. Excerpt from example JSON file retrieved for *ENCODE* experiment ENCSR635OSG.

4.2 Data Transformation

The *Transformer* module takes as input the metadata files output by the previous phase, transforming them into free arbitrary semi-structured key-value pairs, compatible with the Genomic Data Model [22].

Formalization. A *Transformer* is a source-specific method \mathcal{T}_i . When applied to each file in D_i downloaded from a given source i , it produces a file in T_i of $\langle key, value \rangle$ pairs.

Method. The *Transformer* process, exemplified in Algorithm 1, downloads files with an adaptation strategy that depends on their format: (i) hierarchical formats (JSON, XML, or equally expressive) require applying a flattening procedure that creates for each value a pair formed by a key (composed as the concatenation of all JSON/XML elements from the root to the element corresponding to the selected value) and the value itself; (ii) tab-delimited formats (CSV or Excel/Google Sheet) require pivoting tab-delimited columns into rows (which corresponds to creating key-value pairs); (iii) completely unstructured metadata formats, collected from Web pages or other documentation provided by sources, need case-specific formatting. The output of a transformer is a lists of key-value pairs, added to the set T_i . We wrote transformers for the most used formats for origin metadata. Additional ones can be easily added.

Algorithm 1 Transformer Procedure

```
function TRANSFORMATION( $D_i, T_i$ )
  for each  $d \in D_i$  do
    switch  $d$  do
      case  $d$  is hierarchical
         $t \leftarrow flattenPaths(d)$ 
      case  $d$  is tab-delimited
         $t \leftarrow pivot(d)$ 
      case  $d$  is unstructured
         $t \leftarrow manualFormatting(d)$ 
     $T_i \leftarrow T_i + newTransFile(t)$ 
  end for
  return  $T_i$ 
end function
```

Example. The output of data transformation for *ENCODE* is shown in Listing 2; it is obtained by considering as input the portion of the JSON file from Listing 1, which describes the information extracted for a specific item with accession ENCFF429VMY (with two replicates) of experiment ENCSR635OSG. First-level elements are translated directly into $\langle key, value \rangle$ pairs (e.g., $\langle accession, ENCSR635OSG \rangle$); nested elements are flattened (e.g., "name" inside "lab", inside "pi", inside "award" becomes `award_pi_lab_name`, where double underscore `__` is used to separate levels of nesting); arrays are translated into one $\langle key, value \rangle$ pair for each value in the array (e.g., see `file_biological_replicates`); empty arrays are not translated (e.g., "dbxrefs").

Note that several replicates can be associated with each file; in such a case, a progressive naming scheme tracks the replicate to which each $\langle key, value \rangle$ pair relates. In the specific example, the file has two biological replicates, each with five associated key-value pairs (in Listing 2 other pairs are omitted for brevity). All elements in the replicate element with id `4874c170-7124-4822-a058-4bb` are transformed into keys that start with "replicate_1_". Vice versa, elements

in replicate `d42ff80d-67fd-45ee-9159-25a` are transformed into keys that start with "replicate_2_".

```
accession ENCSR635OSG
assembly hg19
award_pi_lab_name michael-snyder
file_accession ENCFF429VMY
file_biological_replicates 1
file_biological_replicates 2
file_file_type bed narrowPeak
replicates_1__id /replicates/4874c170-7124-4822-a058-4bb/
replicates_1__biological_replicate_number 1
replicates_1__library_biosample_donor_age 6
replicates_1__library_biosample_health_status healthy
replicates_1__antibody_lot_id 940739
replicates_2__id /replicates/d42ff80d-67fd-45ee-9159-25a/
replicates_2__biological_replicate_number 2
replicates_2__library_biosample_donor_age 32
replicates_2__library_biosample_health_status healthy
with non-obstructive coronary artery disease
replicates_2__antibody_lot_id 940739
```

Listing 2. Excerpt from example transformed file corresponding to ENCODE file accession ENCFF429VMY.

4.3 Data Cleaning

After the transformation step, a typical key is a long string, e.g., `replicates_1__library_biosample_donor_age`. As this information applies to an ITEM, a much simpler attribute name can be derived, e.g., `donor_1_age`. Such name is later used to map values in the conceptual schema, and is a much simpler key.

The *Cleaner* module applies transformation rules to complex attribute names, so as to simplify them. For illustration purposes, rules are indicated with the notation *antecedent* \Rightarrow *consequent*. The antecedent of rules uses the formalism of *regular expressions*: it recognizes the strings that compose a complex attribute. The consequent, which is an action encoded in the form of *pattern matching replacement strategy*, builds a simpler string. The use of regular expressions brings a simple formalization of cleaning algorithms through language containment and language-recognizing automata. Rules are source-specific, as they depend on the way in which attribute names are encoded at each source; after an initial design, they are applied to each transformed file. Rules may require adjustments when the attribute encoding changes, or new attributes are created. We provide a tool that assists designers in rule creation, ordering, and maintenance.⁵

Informally, rules consist of an antecedent, recognizing an input string, and a consequent, transforming it into a simpler output string. The rule's antecedent is a regular expression matching a sequence of keys; it contains parentheses, which group parts of regular expressions positionally identify the rule's parameters, used in the rule's consequent as numbered capturing groups.⁶ Some parameters are typed, e.g., `[0-9]` denotes a sequence of digits; some keys may be equivalently used, e.g., `(age|sex)` denotes an alternative. The consequent can contain strings of characters or special "dollar" symbols, which positionally refer to the content of the antecedent's variables. The consequent can be empty, in which case no cleaned key is generated for the transformed key, and the corresponding pair is removed.

5. <https://github.com/DEIB-GECO/Metadata-Manager/wiki/Rule-Base-Generator>

6. The replacement strategy specified by a rule is implemented using the `java.util.regex` library (<https://docs.oracle.com/javase/8/docs/api/java/util/regex/package-summary.html>), supporting full regular expressions.

Rule Example. An example of a rule is:

```
replicates(__[0-9]__) library__biosample__(donor)__(age|sex)(.*) => $2$1$3$4.
```

When `replicates__1__library__biosample__donor__age` is considered as the input key, `$2$1$3$4` produces a concatenation of the content of the second variable `donor`, with the first one `__1__`, with the third one `age` and finally with the fourth one (i.e., anything that follows the third parenthesis)—in this case an empty string. As a result, the rule produces the string `donor__1__age`.

Formalization. A *Cleaner* is a source-specific method $C_i = \langle T_i, C_i, RB_i \rangle$. For every transformed metadata file in T_i , it converts the key k of each key-value pair $\langle k, v \rangle$ from its transformed syntax to a cleaned version k' of pair $\langle k', v \rangle$. If k' is empty, a related pair is not produced. Files in C_i contain cleaned key-value pairs and are produced by running the rule engine C_i over T_i using the set of rules RB_i .

Method. The description of the method requires the definition of relations between rules and of rule base.

Definition 1. (*Rule Equivalence, Containment, and Partial Overlap*) Given two rules $r, r' \in RB_i$, their antecedents $r.a$ and $r'.a$, and the corresponding generated languages $\mathcal{L}(r.a)$ and $\mathcal{L}(r'.a)$:

- r is equivalent to r' when $\mathcal{L}(r.a) = \mathcal{L}(r'.a)$;
- r is contained in r' when $\mathcal{L}(r.a) \subset \mathcal{L}(r'.a)$;
- r partially overlaps r' when $\mathcal{L}(r.a) \not\subset \mathcal{L}(r'.a)$, $\mathcal{L}(r'.a) \not\subset \mathcal{L}(r.a)$, and $\mathcal{L}(r.a) \cap \mathcal{L}(r'.a) \neq \emptyset$

Definition 2. (*Rule Base*) The *RB Rule Base* is a list of rules such that rule r precedes rule r' in *RB* if either 1) r is contained in r' , or 2) r partially overlaps r' and the user gives priority to r over r' .

By effect of the above definitions, rules that are more specific precede more general rules. When the intersection of languages recognized by the rules is non-empty, the user can specify the desired order in which the rules should appear in the *RB*. When the intersection is empty, the rules' order in the *RB* corresponds to the order of insertion.

Building a *Cleaner* requires building the Rule Base (Algorithm 2), by calling the function to insert a rule in the right order (Algorithm 3), which is based on the comparison between pairs of rules performed by the function `COMPARE` (Algorithm 4). When the Rule Base is prepared, it is applied to the transformed files, in particular to the keys from the $\langle key, value \rangle$ pairs in T_i (Algorithm 5). After the consolidation of cleaning rules, a rule base can be repeatedly applied to transformed data, until major changes occur at the sources.

Algorithm 2 Rule Base Creation

```
function RBCREATION(RB, SK, AK)
    UK ← AK - SK
    while UK is not empty do
        newRule ← getRuleFromUser()
        if userApprSimul(RB, newRule) then
            RULEINSERTION(RB, newRule)
            matched ← matchAll(RB, UK)
            SK ← SK + matched
            UK ← UK - matched
        end if
    end while
end function
```

Algorithm 2 takes as input *RB*, which stores the information about rules in their order (Definition 1), *SK*, the set of seen keys, and *AK*, the set of all keys retrieved from the files of a given source. It first finds the unseen keys *UK* (those that have not been considered for rule creation yet). Then, until all unseen keys have been considered, the user is asked to insert new rules and approve (or not) the simulated effect of the incremented *RB* on all keys. When the user is satisfied with the results, the rule is actually added to the *RB* and the sets of keys are updated accordingly.

Algorithm 3 Rule addition in Rule Base

```
function RULEINSERTION(RB, newRule)
    for r in RB do
        res ← COMPARE(newRule, r)
        if res is EQUIVALENT then
            if userPref(newRule, r) = newRule then
                replaceRule(newRule, RB, indexOf(r))
            end if
            return RB
        else if res is CONTAINED then
            addRule(newRule, RB, indexOf(r))
            return RB
        else if res is PARTIALLY_OVERLAPS then
            if userPriority(newRule, r) = newRule then
                addRule(newRule, RB, indexOf(r))
            end if
            return RB
        end if
    end for
    addRule(newRule, RB, RB.size)
    return RB
end function
```

Adding a new rule to the Rule Base means inserting it in the correct position with respect to the order defined in Definition 2. This is accomplished by Algorithm 3, which iterates over the *RB* list and, based on the comparison between each pre-existing rule with the one to be added, determines the insertion position (in an "insertion sort" manner).

Algorithm 4 Order comparison between rules

```
function COMPARE(r, r')
    Ar ← NFA2DFA(RegEx2NFA(r.a))
    A'r ← NFA2DFA(RegEx2NFA(r'.a))
    if L(Ar) = L(A'r) then
        return EQUIVALENT
    else if L(Ar) ⊂ L(A'r) then
        return CONTAINED
    else if L(Ar) ⊄ L(A'r) ∧ L(Ar ∩ A'r) ≠ ∅ then
        return PARTIALLY_OVERLAPS
    end if
end function
```

Algorithm 5 Application of Rule Base to keys

```
function CLEANER(RB, Ti)
    Ci ← []
    for each (key, value) ∈ Ti do
        newKey ← matchFirst(key, RB)
        if nonEmpty(newKey) then
            add(Ci, (newKey, value))
        end if
    end for
    return Ci
end function
```

Comparing rules means evaluating the containment relationship between the languages generated by their an-

tecedents, as described by Algorithm 4. Several procedures exist to convert regular expressions into equivalent Non-deterministic Finite Automata (NFA); we use the Brics Java library [23] for automata implementations, which is based on Thompson's construction algorithm [24]. Then, NFA need to be converted into equivalent Deterministic Finite State Automata (DFA) \mathcal{A}_r and $\mathcal{A}_{r'}$ —this can be done with the Rabin-Scott powerset construction [25]. Later, the two languages are checked for equivalence, containment and partial overlapping (by using the automaton constructed from the cross-product of states that accepts the intersection of the languages). Algorithm 5 describes the *Cleaner* as application of the Rule Base to the input dataset; rules are applied in the order in which they appear in the Rule Base.

Example. Table 2 illustrates a cleaning process for *ENCODE*. It assumes an initial set of transformed keys from T_i ; for each key, the user produces cleaning rules, driven by Algorithm 2. Eventually, the method produces a rule base of 7 rules; their application to keys in T_i produces the set of cleaned keys in C_i .

TABLE 2
Example of cleaning process.

| Transformed keys in T_i | |
|--|----------------|
| replicates_1_library_biosample_donor_age | 32 |
| replicates_1_library_biosample_donor_age_units | year |
| replicates_1_library_biosample_donor_sex | male |
| replicates_2_library_biosample_donor_age | 4 |
| replicates_2_library_biosample_donor_age_units | year |
| replicates_2_library_biosample_donor_sex | female |
| replicates_1_library_biosample_sex | male |
| replicates_1_library_biosample_biosample_type | tissue |
| replicates_1_library_biosample_health_status | healthy, CAD |
| file_biological_replicates | 1 |
| file_technical_replicates | 1_1 |
| file_assembly | GRCh38 |
| file_file_type | bed narrowPeak |
| replicates_1_biological_replicate_number | 1 |
| replicates_1_technical_replicate_number | 1 |
| replicates_2_biological_replicate_number | 2 |
| replicates_2_technical_replicate_number | 1 |
| assembly | hg19 |

↓

| RuleBase RB_i | |
|--|----------------|
| (1) replicates([0-9]_)library_biosample_(donor)(age sex)(.*) | ⇒ \$2\$1\$3\$4 |
| (2) replicates_[0-9]_library_biosample_.sex.* | ⇒ |
| (3) replicates([0-9]_)library_(biosample)_?(biosample)_?(.*) | ⇒ \$2\$1\$4 |
| (4) file_(biological technical)_replicates | ⇒ |
| (5) (file_)(file_)?(.*). | ⇒ \$1\$3 |
| (6) (replicate)s([0-9]_)(.*) | ⇒ \$1\$2\$3 |
| (7) assembly | ⇒ |

↓

| Cleaned keys in C_i | |
|---|----------------|
| donor_1_age | 32 |
| donor_1_age_units | year |
| donor_1_sex | male |
| donor_2_age | 4 |
| donor_2_age_units | year |
| donor_2_sex | female |
| biosample_1_type | tissue |
| biosample_1_health_status | healthy, CAD |
| file_assembly | GRCh38 |
| file_type | bed narrowPeak |
| replicate_1_biological_replicate_number | 1 |
| replicate_1_technical_replicate_number | 1 |
| replicate_2_biological_replicate_number | 2 |
| replicate_2_technical_replicate_number | 1 |

For instance, rule (2) deletes the key: `replicates_1_library_biosample_.sex`. Rule (3), applied to the key: `replicates_1_library_biosample_biosample_type`, dictates that the key must be rewritten by concatenating the content of the second parenthesis (i.e., `biosample`) with the content of the first (i.e., `1`), and with the content of the

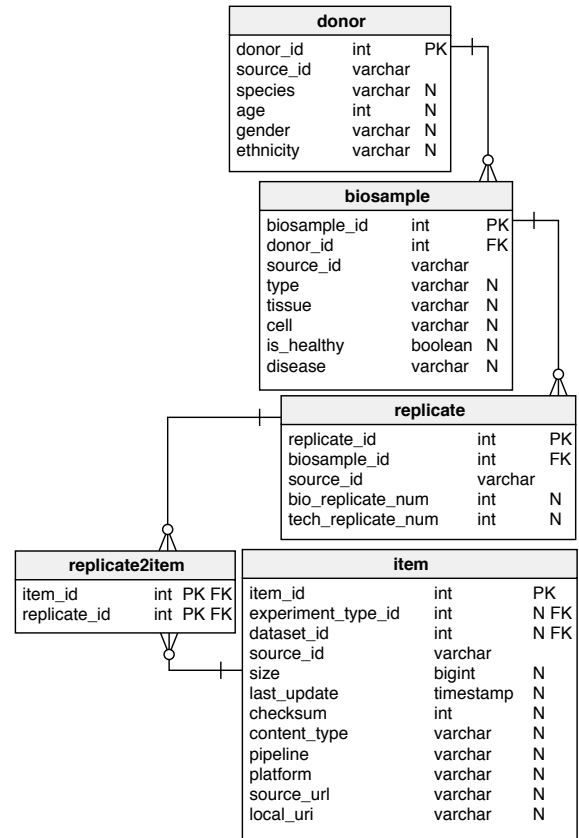


Fig. 5. Logical schema of the GCM biological view.

fourth (i.e., `type`), obtaining at the end `biosample__1__type`.

5 DATA INTEGRATION

The META-BASE data integration process consists of three phases. During data mapping, described in Section 5.1, cleaned metadata is mapped into a global relational schema that embodies the conceptual schema presented in Section 2. Data mapping is a simple syntactic transformation; the following phase of value normalization and enrichment, described in Section 5.2, produces homogenized data equipped with appropriate ontological term labels, references, hyponyms, hypernyms and synonyms. Finally, the integrity constraint checker, discussed in Section 5.3, provides methods for specifying and enforcing integrity constraints that describe legal values in the META-BASE repository. For a high-level workflow of the data integration process refer again to Fig. 3.

5.1 Data Mapping

The *Mapper* module is in charge of the integration at the schema-level of a set of cleaned keys produced for each source. The method extends the work proposed in [8], where we first introduced local-to-global mappings using a classical Datalog [26] syntax. The current *Mapper* is part of a broader integration workflow in which metadata is made available as lists of $\langle key, value \rangle$ pairs; mapping rules build relational rows from such pairs.

The global relational schema \mathcal{G} is obtained as straightforward mapping from the conceptual schema in Fig. 1.

It contains the central entity table ITEM, a set of entity tables DONOR, BIOSAMPLE, REPLICATE, PROJECT, CASE, DATASET, EXPERIMENTTYPE, which model as well 1:N relationships, and two relationship tables ITEM2REPLICATE and ITEM2CASE, which model the two N:N relationships. Fig. 5 shows the logical schema of the biological view, where *PK* denotes attributes forming the table's primary key, *FK* denotes foreign keys, *N* denotes nullable attributes, multiplicity in the edge denotes a many mapping, a circle on the edge denotes an optional mapping, and a cut on the edge denotes a mandatory mapping.

Every source is represented by a set of cleaned files C_i , each of which contains a set of $\langle key, value \rangle$ pairs, where the *keys* are produced by the *Cleaning* phase. Mapping rules assemble several values extracted from the key-value pairs into rows of the relations in \mathcal{G} . Their format recalls deductive rules: each table of \mathcal{G} corresponds to several rules for each source, whose head is a predicate named as the table and with the same arity as the table's grade; the body lists several attribute-value pairs such that attribute names are matched to cleaned keys of files in C_i . The semantics of mapping rules is also similar to that of deductive rules: if all the attribute names of the body are matched to keys in C_i (in deductive terms they *unify*), then the values corresponding to those keys are assembled by the rule into relational rows.

TABLE 3
Example syntactic transformations for mapping rules.

| |
|--|
| Conc(s_1, s_2, c): concatenates s_1 and s_2 using c as separation string |
| Alt(s_1, s_2): outputs s_1 if present and not null, else s_2 |
| Rem(s_1, s_2): removes the occurrences of string s_2 from s_1 |
| Sub(s_1, s_2, s_3): substitutes occurrences of s_2 in s_1 with the new s_3 |
| Eq(s, p): outputs true when s is equal to p , else false |
| ATD(a): converts a , a number followed by space and unit of measurement, into the correspondent number of days |
| LCase(s): converts string s into its lower case version |
| Int(n): casts number n to its correspondent Integer format |
| Id(...): generates synthetic <i>id</i> for faster indexing of table t from specified arguments |

It is possible to apply to the values a set of pre-defined syntactic transformations (*SynTr*), defined in Table 3, which can be freely composed in mapping rules; for example, to put into lowercase letters two values that have been first concatenated with a space, the expression $LCase(Conc(value_1, value_2, " "))$ can be used to generate a value for a specific position of a row. Transformations can be easily extended.

Formalization. A *Mapper* is a source-specific method $\mathcal{M}_i = \langle C_i, MB_i \rangle$. For every cleaned metadata file f in C_i , it assembles several values v present in the pairs $\langle k, v \rangle$ of f into rows of the tables of \mathcal{G} . Rows are produced by running the rule engine \mathcal{M}_i over C_i using the mapping rules in MB_i . A *mapping rule* is a declarative rule of the form: $ENTITY(SynTr(v_1), \dots, SynTr(v_i), \dots, SynTr(v_N)) \rightsquigarrow \{ \langle k_1, v_1 \rangle, \dots, \langle k_i, v_i \rangle, \dots, \langle k_N, v_N \rangle \} \subseteq f$. The left-hand side (LHS) and the right-hand side (RHS) of the rule are w.r.t. the \rightsquigarrow symbol; every v_i in the LHS of the rule also appears in the RHS of the rule (i.e., rule evaluations are *finite*), in a positive form (i.e., rules are *safe*).

Method. Once mapping rules are fully specified, the method consists simply in applying the rules to each file f in C_i of data source i , in arbitrary order. Note that every

file associated with a data source, as produced by the cleaning method, may have several versions for the same key, numbered from 1 to n_f (e.g., $biosample_1_type$, $biosample_2_type$); each rule is applied for every version, and associates it with a distinct row. When a version is present (e.g., in rules for DONOR, BIOSAMPLE and REPLICATE of Table 4), we denote it by generically naming the keys in the rule's RHS using j , and then generating a rule for each value of j . For each v_i in the rule's LHS, if the corresponding $\langle k_i, v_i \rangle$ in the rule's RHS exists in f , then we add $SynTr(v_i)$ to the result, i.e., a tuple in the ENTITY table specified in the rule's LHS.

Example. Table 4 illustrates all the rules required to build the relational schema shown in Fig. 5 for the ENCODE data source. Note that oid_t is the notation used for the *ObjectIdentifier* of table t , which is a unique accession identifier retrieved from the source.

TABLE 4
Mapping rules for biological view of ENCODE source.

| | |
|--|---|
| $C_{ENC} = \{C_i i = ENCODE\}, \forall f \in C_{ENC}, j \leq n_f$ | |
| DONOR($Id(oid_D), oid_D, v_1, ATD(Conc(v_2, v_3, " ")), v_4, v_5$) \rightsquigarrow | $\{ (donor_j_accession, oid_D), (donor_j_organism, v_1), (donor_j_age, v_2), (donor_j_age_units, v_3), (donor_j_sex, v_4), (donor_j_ethnicity, v_5) \} \subseteq f$ |
| BIOSAMPLE($Id(oid_B), Id(oid_D), oid_B, "tissue", v_2, NULL, Eq(v_3, "healthy"), v_3$) \rightsquigarrow | $\{ (biosample_j_accession, oid_B), (donor_j_accession, oid_D), (biosample_j_type, "tissue"), (biosample_j_term_name, v_2), (biosample_j_health_status, v_3) \} \subseteq f$ |
| BIOSAMPLE($Id(oid_B), Id(oid_D), oid_B, "cell\ line", NULL, v_2, Eq(v_3, "healthy"), v_3$) \rightsquigarrow | $\{ (biosample_j_accession, oid_B), (donor_j_accession, oid_D), (biosample_j_type, "cell"), (biosample_j_term_name, v_2), (biosample_j_health_status, v_3) \} \subseteq f$ |
| REPLICATE($Id(oid_R), oid_R, oid_B, v_1, v_2$) \rightsquigarrow | $\{ (replicate_j_uid, oid_R), (biosample_j_accession, oid_B), (replicate_j_bio_rep_num, v_1), (replicate_j_tech_rep_num, v_2) \} \subseteq f$ |
| ITEM($Id(oid_I), Id(v_1, v_2, v_3), Id(oid_DS), oid_I, v_4, v_5, v_6, v_7, v_8, Conc("www.encodeproject.org", v_9, "/"), Conc("www.gmql.eu...", oid_I, "/")$) \rightsquigarrow | $\{ (assay_term_name, v_1), (target_investigated_as, v_2), (target_label, v_3), (dataset_name, oid_DS), (file_accession, oid_I), (file_size, v_4), (file_date_created, v_5), (file_md5sum, v_6), (file_pipeline, v_7), (file_platform, v_8), (file_href, v_9) \} \subseteq f$ |
| ITEM2REPLICATE($Id(oid_I), Id(oid_R)$) \rightsquigarrow | $\{ (file_accession, oid_I), (replicate_j_uid, oid_R) \} \subseteq f$ |

5.2 Data Normalization and Enrichment

During this step, specific values of the global schema are associated with controlled terms, lists of synonyms and hypernyms, and external references to specialized ontologies. We consider nine *semantically enrichable* attributes of the global schema: *Technique*, *Feature* and *Target* of experiment types; *Disease*, *Tissue* and *Cell* of bio samples; *Ethnicity* and *Species* of donors; *Platform* of items.

The adoption of a specific knowledge base for each semantically enrichable attribute provides us with *value normalization*, as we consider the values of reference knowledge bases as a restricted vocabulary. Using external knowledge bases (rather than creating a new one) is essential in the

biomedical domain, where specialized ontologies are already available and their use boosts interoperability.

This process is supervised and requires a preliminary selection of the most suitable ontologies to describe each semantically enrichable attribute of the global schema.

5.2.1 Ontology Selection

The choice of attribute-specific ontologies took into account the rules for selecting a bio-ontology given in [27]. Among others available, we used four different services to evaluate the best ontologies for nine ontological attributes from GCM. These are: (a) BioPortal [28], (b) Ontology Recommender [29], (c) Ontology Lookup Service (OLS, [20]), and (d) Zooma.⁷ For each semantically enrichable attribute, we searched all values using the four services, and computed the best score for recommended ontologies.⁸ Finally, for each pair attribute-ontology, we considered both the best matching scores and the coverage (number of values of a given attribute that were successfully annotated, i.e., matched to an ontological term).

The results of our selection are shown in Table 5; for each considered attribute, we indicate the preferred ontology and three normalized indicators. COVERAGE indicates the percentage of attribute values that are found in the ontologies. SCORE is an average matching score of all the annotated attribute values weighted by ontology acceptance. SUITABILITY is a measure of how much an ontology set is adequate for an attribute. Note that a second preferred ontology is added when the first one did not reach 0.85 coverage. In this case, indicators refer to the union of the ontologies. Details on the procedure are given in [30].

TABLE 5
Choice of reference ontologies for semantically enrichable attributes.

| Attribute | Pref. ontologies | Coverage | Score | Suitability |
|-----------|------------------|----------|-------|-------------|
| Technique | OBI, EFO | 0.857 | 0.486 | 0.490 |
| Feature | NCIT | 1.000 | 0.854 | 0.893 |
| Target | OGG | 0.950 | 0.747 | 0.948 |
| Disease | NCIT | 0.978 | 0.784 | 0.802 |
| Tissue | UBERON | 0.957 | 0.753 | 0.937 |
| Cell | EFO, CL | 0.953 | 0.644 | 0.577 |
| Platform | NCIT | 1.000 | 0.909 | 0.950 |
| Ethnicity | NCIT | 0.962 | 0.907 | 0.912 |
| Species | NCBITaxon | 1.000 | 0.667 | 1.000 |

5.2.2 Process

The *Normalizer* is supported by an interactive tool that: 1) calls external services to annotate values with concepts from controlled vocabularies or dedicated ontologies; 2) asks for user feedback when annotations have a low matching score; users can either accept one of the proposed solutions, or manually specify new annotations.

The result of the normalization is contained within the relational database \mathcal{K} , called *Local Knowledge Base*, illustrated in Fig. 6, populated from ontologies and referenced from the global schema \mathcal{G} . Specifically, we maintain the tables:

7. <https://www.ebi.ac.uk/spot/zooma/>

8. Recommender provides numerical scores, Zooma provides tags for indicating the annotation quality. For BioPortal and OLS we computed a score by considering the number of words with exact match in each ontology.

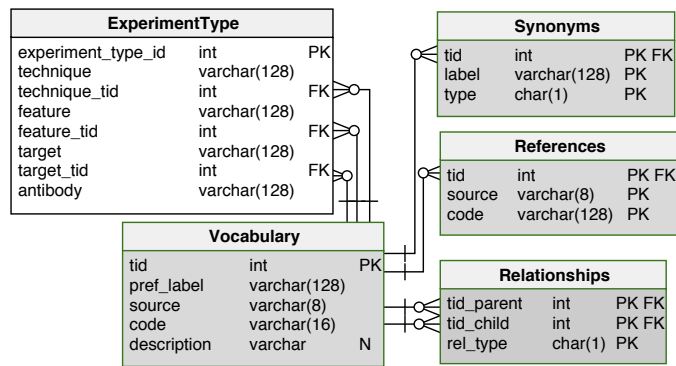


Fig. 6. Relational schema of the Local Knowledge Base \mathcal{K} , including links to attributes of EXPERIMENTTYPE from the global schema \mathcal{G} .

- 1) VOCABULARY: contains the term identifier and the term preferred label, in addition to the ontology providing the label, the code used for the label in that ontology and an optional description.
- 2) REFERENCES: for a given term, contains references to equivalent labels extracted from other ontologies (in the form of a pair $\langle Source, Code \rangle$).
- 3) SYNONYMS: contains other labels that can be used as synonyms of the preferred label in the chosen ontology.
- 4) RELATIONSHIPS: contains ontological hierarchical relationships between terms and the type of the relationships (either generalization or containment).

The system provides the unfolding of the hierarchies as an internal materialized view over the table RELATIONSHIPS, used for faster query processing.

Formalization. The *Normalizer* is a source-independent method $\mathcal{N} = \langle A, \mathcal{O} \rangle$. A is the set of semantically enrichable attributes of the global schema \mathcal{G} . For each attribute a in A and each possible value of a , \mathcal{N} generates the corresponding entries in the Local Knowledge Base \mathcal{K} , extracted from the preferred ontologies of a in \mathcal{O} , i.e., the set of reference ontologies defined by the process described in Section 5.2.1.

Method. Value normalization and enrichment is a supervised procedure illustrated in Fig. 7. The workflow is executed for all values of semantically enrichable attributes, and consists of two parts: 1) For each such value, the system initially looks for a suitable term in the vocabulary of the Local Knowledge Base; if a match is available, and the term was already annotated in the past, the procedure is completed. When the match is successful but annotations are lacking, a user's feedback is requested. 2) Terms that do not match with the vocabulary, or whose annotations are not approved by the user, are then searched within the specific ontologies associated with the attribute, as defined in Section 5.2.1. If matches are of high confidence (i.e., matching score above a given threshold), the procedure is completed; if the confidence is low, user feedback is requested. When feedback is negative or there is no match, users are asked to manually provide a new vocabulary term.

Example. Fig. 8 shows a tuple of the BIOSAMPLE global table. Solid line nodes include normalized attribute values; dashed line nodes represent some of the synonyms; dotted line nodes represent hierarchies, labeled by the relevant ontology (only a small subset is represented for brevity).

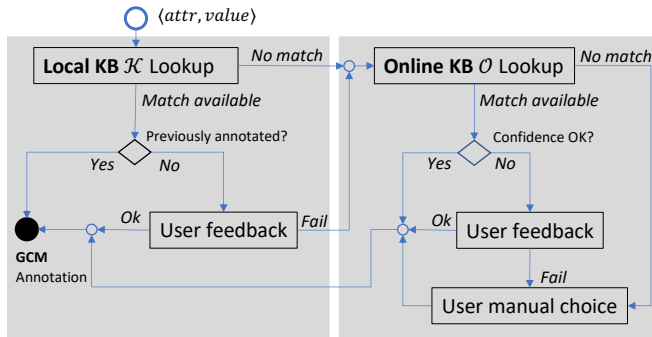


Fig. 7. Iterative supervised normalization and enrichment procedure.

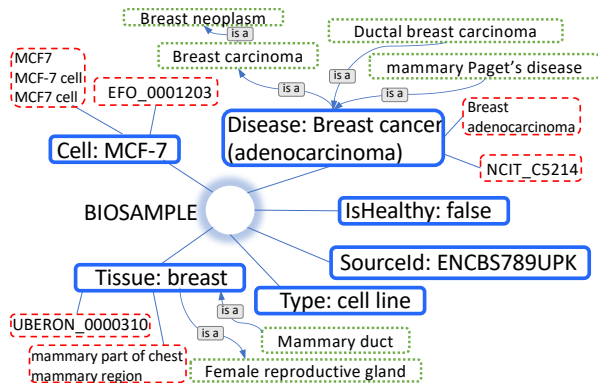


Fig. 8. Normalization and enrichment of a BIOSAMPLE tuple.

5.3 Integrity Checker

At the end of the integration process, we introduce integrity constraints, which define dependencies between values of the global schema \mathcal{G} . Preliminary versions of dependencies, called contextual and dependent features, were introduced in [8]. We consider pairs of attributes $(A_S \in R_S$ and $A_E \in R_E)$, where R_S and R_E denote the starting and ending tables in the \mathcal{G} global schema, connected by a join path in \mathcal{G} . Given that \mathcal{G} is an acyclic schema, there is just one join path between any two tables in \mathcal{G} .

Definition 3. A dependency rule between attributes $R_S.A_S$ and $R_E.A_E$ of \mathcal{G} is an expression of the form: $\text{Boolean}(R_S.A_S) \rightarrow \text{Boolean}(R_E.A_E)$, where $\text{Boolean}(A)$ is a Boolean expression over an attribute A of \mathcal{G} . The interpretation of the dependency rule is that: **when** the Boolean expression in the left part of the rule is true for a value $v_S \in A_S$, **if** there exists one value $v_E \in A_E$ such that $\langle v_S, v_E \rangle$ are connected by the join path between R_S and R_E , **then** the Boolean expression on the right part of the rule must be true for $v_E \in A_E$. Boolean expressions include predicates IS NULL/IS NOT NULL as special cases.

Dependencies can be defined during the lifetime of the META-BASE repository; they are manually defined and their identification is not assisted by a tool, even if we are evaluating to aid the production of such rules using an Association Rules Mining approach inspired by [31]. Table 6 shows some examples of dependency rules. For example, the first rule indicates that if the *Species* of a DONOR is “Homo sapiens” and the donor is connected to a DATASET through the only possible path in \mathcal{G} , then the *Assembly* of the dataset must be one of “hg19”, “hg38”, or “GRCh38”. Dependency rules allow including in the GCM

relevant attributes that are not common to all data types. For example, attributes *Target* and *Antibody* of EXPERIMENTTYPE are of great interest in ChIP-seq experiments, but are not significant in other experiments; a rule can specify that when *Technique* is not “ChIP-seq”, then these attributes are null.

TABLE 6

Examples of dependency rules, including a description of the join path connecting the two attributes used in the left and right parts of the rule.

| |
|---|
| $\langle e_S, e_E \rangle$ in (DONOR·BIOSAMPLE·REPLICATE·ITEM·DATASET) |
| $e_S.Species = \text{“Homo sapiens”} \rightarrow e_E.Assembly \in \{hg19, hg38, GRCh38\}$ |
| $\langle e_S, e_E \rangle$ in (DONOR·BIOSAMPLE) |
| $e_S.Gender = \text{“Male”} \rightarrow e_E.Disease \neq \text{“Ovarian cancer”}$ |
| $e_E.Tissue \neq \text{“Uterus”}$ |
| $e_S.Gender = \text{“Female”} \rightarrow e_E.Disease \neq \text{“Prostate cancer”}$ |
| $\langle e_S, e_E \rangle$ in (PROJECT·CASE·ITEM) |
| $e_S.ProgramName = \text{“ENCODE”} \rightarrow e_E.SourceId = \text{“ENCF.*”}$ |
| $e_S.ProgramName = \text{“TCGA”} \rightarrow e_E.SourceId =$ |
| $\text{“}^\wedge[0-9a-z]\{8\}-\{[0-9a-z]\{4\}\}\{3\}[0-9a-z]\{12\}$”$ |
| $e_S.ProgramName = \text{“ENCODE”} \rightarrow e_E.SourceUrl$ is not null |
| $e_S.ProgramName = \text{“TCGA”} \rightarrow e_E.SourceUrl$ is not null |
| $\langle e_S, e_E \rangle$ in (BIOSAMPLE) |
| $e_S.Type = \text{“tissue”} \rightarrow e_E.Tissue$ is not null |
| $e_S.Type = \text{“cell line”} \rightarrow e_E.Cell$ is not null |
| $\langle e_S, e_E \rangle$ in (PROJECT·CASE·ITEM·DATASET) |
| $e_S.ProgramName = \text{“ENCODE”} \rightarrow e_E.Name = \text{“*ENCODE.*”}$ |
| $e_S.ProgramName = \text{“Roadmap Epigenomics”} \rightarrow e_E.Name =$ |
| $\text{“*ROADMAP_EPIGENOMICS.*”}$ |

6 VALIDATION

We verified the META-BASE architecture from three perspectives: 1) *lossless integration*: the process does not miss information available in the original sources; 2) *semantic enrichment*: semantic information complementing original metadata helps in identifying relevant and additional items; 3) *use evaluation*: we measure the suitability of the repository and of its user interfaces for the bioinformatics community.

Lossless integration. As the generation of all possible queries and their manual check is not feasible, we performed the evaluation on a restricted number of meaningful example queries that show the effectiveness of our approach. This evaluation is hard due to a number of issues, including among others the following ones: i) query interfaces at sources are different from ours (e.g., free-text search vs. attribute-based search); ii) sources assign metadata to different entities (e.g., experiments vs. data files included in experiments)

In our work, we carefully considered all these issues: we studied the source query mechanisms at our best and favored an exact-match strategy for the benefit of comparison. Table 7 reports the number of items resulting from seven exact-match queries as found either in the META-BASE repository or in individual sources. Note that the queries “H3K27me3”, “fat”, and “Illumina Genome Analyzer II” returned a number of matches in our system equal to the sum of matches in the integrated sources. However, Table 7 shows also some unavoidable difficulties, such as *semantic mismatch* (i.e., items are *false positives* when they match the searched string but their attribute’s semantic is different from the one intended by the search). Indeed, the query “MCF-7” suffers from a problem with ENCODE matches: out of 6,442, only 1,322 (the ones found by our system for the same source) actually refer to the value

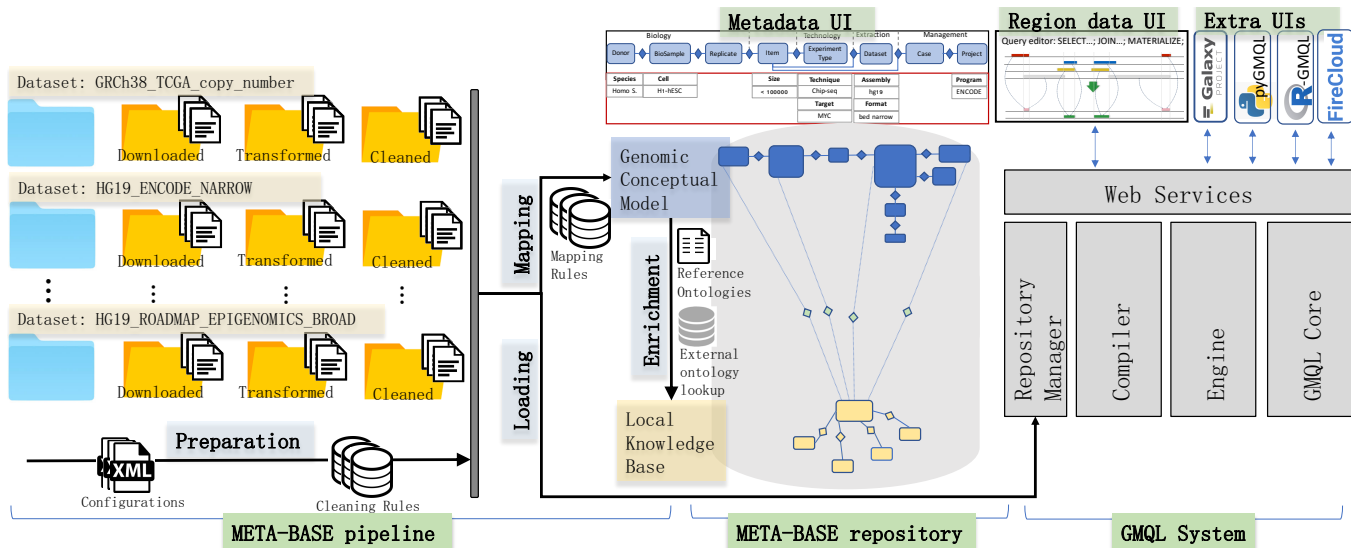


Fig. 9. Overall architecture for genomic data processing.

describing the cell line of the item, while the remaining ones match with alternative information, regarding *possible controls*, *revoked files*, or *summary* of the experiment. Likewise, in *ENCODE* the query “RNA-seq” matches 5,503 items only because this information is contained in *related series* attributes, thus not describing the technique of the item. Similarly, “breast” and “breast cancer” (to a lesser extent) present a considerable number of false positives in *GDC*: in the first case 1,509 items are wrongly matched due to information regarding family members (e.g., tagged with XML key “clinical_patient.family_history_cancer_type”), or information that do not directly describe the sample.

TABLE 7

Comparison of # items from exact-match queries in META-BASE vs. other sources. (ENC=*ENCODE*, Cis=*Cistrome*, 1KG=*1000 Genomes*).

| Attribute | Query | META | ENC | GDC | REP | Cis | 1KG |
|-----------|-----------------------------|--------|-------|--------|-----|-----|-----|
| Target | H3K27me3 | 1,990 | 814 | - | 381 | 795 | - |
| Cell | MCF-7 | 1,447 | 6,442 | - | - | 125 | - |
| Tissue | fat | 57 | - | - | 57 | - | - |
| Tissue | breast | 23,729 | 92 | 24,788 | 94 | 264 | - |
| Disease | breast cancer | 114 | - | 45 | - | 163 | - |
| Platform | Illumina Genome Analyzer II | 981 | 723 | - | - | - | 258 |
| Technique | RNA-seq | 56,047 | 5,503 | 55,650 | 399 | - | - |

Semantic enrichment. This aspect has been validated by six biology experts, who have been asked to evaluate a random set of 200 matches achieved by our supervised procedure between metadata values and ontological controlled terms (either *preferred label* or *synonym*) equipped with their descriptions. Overall, 92% of results were considered “exact”, “good” or “acceptable”; the process is detailed in [30].

Use evaluation. The usefulness of the META-BASE content, accessed using *GenoSurf* [14], was evaluated by asking 40 informed users (i.e., computational biologists, bioinformaticians and data scientists with interest in life sciences) to complete a questionnaire of 10 composite queries and to fill a form reporting on their overall experience with the platform. Detailed results are shown in [15], overall indicating positive feedbacks.

7 OVERALL ARCHITECTURE

META-BASE is part of a broad architecture, whose main purpose is providing a cloud-based environment for genomic data processing. The overall system architecture is presented in Fig. 9. In the left part of the figure we show the META-BASE pipeline discussed in Sections 4 and 5; the whole pipeline is configured using parameters provided as a single XML configuration file. Each dataset (on the left) is progressively downloaded, transformed and cleaned. The data mapping method transforms cleaned attribute-value pairs into the global database \mathcal{G} . The normalization and enrichment method adds references from the semantically enrichable attributes to the Local Knowledge Base \mathcal{K} , which is implemented by relational tables. Interactive access to the META-BASE repository is provided by a user-friendly interface (not discussed in this paper), that uses the acyclic structure of the global schema to support simple conjunctive queries at the center of Fig. 9), in a style that is similar to *DeepBlue’s* query interface [32].

As shown in the right part of Fig. 9, the META-BASE repository can also be queried using the *GMQL System* [16], which supports integrated data management on the cloud; the system is accessed through *Web Services* as a common point of access from a variety of interfaces, including a visual user interface, programmatic interfaces for Python [33] and R/Bioconductor, and workflow-based interfaces for *Galaxy* and *FireCloud*. The implementation is executed using the *Apache Spark* engine, deployed either on a single server or a cloud-based system.

8 RELATED WORK

Genomic repositories are growing in number, diversity and complexity (see last *Nucleic Acid Research* yearly report [34]). In this context, data integration challenges in the *omics* domain are many, as reviewed in [35], [36]. There have been efforts to integrate multiple projects in single initiatives or portals, proposing a unifying data model or strategy to offer integrated access and management of genomic data. Federated access of several databases is offered

by the BioMart [37] project, while Pathway Commons [38] collects pathways from different databases into a shared repository, and ExPASy [39] links many resources in life sciences. Examples among consortia include: large epigenomic players – e.g., ENCODE Data Coordination Center [40], the International Human Epigenome Consortium [41], and DeepBlue [32] of the Blueprint Consortium; cancer-related consortia – e.g., GDC (data model: <https://gdc.cancer.gov/developers/gdc-data-model>) and ICGC [42]. All such actors do not provide so far neither models nor integration frameworks that are general enough to cover aspects falling outside the specific focus of their scope. Differently, the META-BASE approach is independent from specific sub-branches of genomics and can be applied to a large number of heterogeneous sources by any integration designer, in *a-posteriori* fashion, i.e., without having to follow any guidelines in the preliminary production of metadata.

Instead of addressing the problem of integration, some works promote methodologies to improve the process of *metadata authoring* (i.e., preliminary preparation and submission) or the manual curation of metadata. Among these: CEDAR [43], a system for development, evaluation, use, and refinement of genomics and biomedical metadata, using ontology-based recommendations from BioPortal [28]; BioSchemas.org [44], which applies schemata to online resources making them easily searchable; DNAdigest [45], promoting efficient sharing of human genomic datasets; DATS [46] boosting datasets' discoverability.

Many works in the literature use conceptual models (CMs) in the biomedical field [47], but they employ CMs expressive power to understand and explain biological entities and their interactions [48], [49], [50], [51]. Instead, we propose in addition an architecture that uses a conceptual model for driving data integration.

The process of semantic enrichment of metadata, including the choice of appropriate ontologies [52] has been tackled both with source-independent methods [53], [54] and with source-specific ones (for GEO [55], for ENCODE [56]).

9 DISCUSSION AND CONCLUSIONS

Genomic metadata integration is a complex process. We have designed our solution by breaking the process into several tasks and by associating powerful abstractions to each task. The *Cleaner*, *Mapper* and *Checker* modules contain new methods, all based on the interplay of different kinds of rules. The *Cleaner* rules are inspired to grammar-based transformations (where order matters), the *Mapper* rules are expressed in a Datalog-like formalism (order independent), the *Checker* rules are expressed as logical dependencies (order independent). Each rule-based method is provided with a formal description. The iterative process used in normalization and enrichment, where the designer's feedback is needed in order to validate or suggest annotations, has never been applied to genomic metadata. The significance of our approach stands in providing a single framework where the interplay of the three kinds of rules and of the effective interaction with the designer drives the whole process.

The META-BASE repository currently includes datasets from the ENCODE project (≈ 21 million metadata $\langle key, value \rangle$ pairs in 26,111 items from 4 datasets), GDC

(TCGA program, with $\approx 18M$ pairs, $\approx 100K$ items, 7 datasets), Roadmap Epigenomics ($\approx 200K$ pairs, $\approx 3K$ items, 6 datasets), Cistrome ($\approx 80K$ pairs, $\approx 6K$ items, 2 datasets), GENCODE and RefSeq annotations ($\approx 1.3K$ pairs, 105 items, 4 datasets), and TADs from GEO (272 pairs, 14 items, 2 datasets). In total, our framework has imported $\approx 40M \langle key, value \rangle$ pairs, which correspond to $\approx 3.1K$ distinct keys. The addition of new sources to the repository is ongoing; our current commitment is to extend the data integration process by adding relevant sources one-by-one, in an incremental fashion. The resulting META-BASE repository is an important resource for supporting biological and clinical research. META-BASE can be easily re-instantiated in other scientific database contexts where the problem of origin heterogeneous sources also applies.

ACKNOWLEDGMENTS

This research is funded by the ERC Advanced Grant 693174 GeCo (data-driven Genomic Computing). The authors would like to thank the contributions of Alessandro Campi and of the master students Andrea Colombo, Federico Gatti, Riccardo Mologni, and Jorge Ignacio Vera Pena.

REFERENCES

- [1] Consortium ENCODE, "An integrated encyclopedia of DNA elements in the human genome," *Nature*, vol. 489, no. 7414, pp. 57–74, 2012.
- [2] J. N. Weinstein *et al.*, "The Cancer Genome Atlas pan-cancer analysis project," *Nature Genetics*, vol. 45, no. 10, pp. 1113–1120, 2013.
- [3] M. A. Jensen *et al.*, "The NCI Genomic Data Commons as an engine for precision medicine," *Blood*, vol. 130, no. 4, pp. 453–459, 2017.
- [4] A. Kundaje *et al.*, "Integrative analysis of 111 reference human epigenomes," *Nature*, vol. 518, no. 7539, pp. 317–330, 2015.
- [5] 1000 Genomes Project Consortium, "A global reference for human genetic variation," *Nature*, vol. 526, no. 7571, p. 68, 2015.
- [6] A. Palacio León and O. Pastor López, "From big data to smart data: A genomic information systems perspective," in *Proc. 12th Int. Conf. RCIS. IEEE*, 2018, pp. 1–11.
- [7] R. S. Gonçalves and M. A. Musen, "The variable quality of metadata about biological samples used in biomedical experiments," *Sci. Data*, vol. 6, p. 190021, 2019.
- [8] a. Bernasconi *et al.*, "Conceptual modeling for genomics: Building an integrated repository of open data," in *Proc. 36th Int. Conf. ER. Springer*, 2017, pp. 325–339.
- [9] A. Frankish *et al.*, "GENCODE reference annotation for the human and mouse genomes," *Nucleic Acids Res.*, vol. 47, no. D1, pp. D766–D773, 2019.
- [10] N. A. O'Leary *et al.*, "Reference sequence (RefSeq) database at NCBI: current status, taxonomic expansion, and functional annotation," *Nucleic Acids Res.*, vol. 44, no. D1, pp. D733–D745, 2016.
- [11] S. S. Rao *et al.*, "A 3D map of the human genome at kilobase resolution reveals principles of chromatin looping," *Cell*, vol. 159, no. 7, pp. 1665–1680, 2014.
- [12] T. Barrett *et al.*, "NCBI GEO: archive for functional genomics data sets—update," *Nucleic Acids Res.*, vol. 41, no. D1, pp. D991–D995, 2012.
- [13] R. Zheng *et al.*, "Cistrome Data Browser: expanded datasets and new tools for gene regulatory analysis," *Nucleic Acids Res.*, vol. 47, no. D1, pp. D729–D735, 2018.
- [14] A. Canakoglu *et al.*, "GenoSurf: metadata driven semantic search system for integrated genomic datasets," *Database*, vol. 2019, 12 2019, baz132.
- [15] A. Bernasconi *et al.*, "Exploiting conceptual modeling for searching genomic metadata: A quantitative and qualitative empirical study," in *Adv. in Conceptual Modeling. Springer*, 2019, pp. 83–94.
- [16] M. Masseroli *et al.*, "Processing of big heterogeneous genomic datasets for tertiary analysis of Next Generation Sequencing data," *Bioinformatics*, vol. 35, no. 5, pp. 729–736, 2019.

- [17] —, “GenoMetric Query Language: a novel approach to large-scale genomic data management,” *Bioinformatics*, vol. 31, no. 12, pp. 1881–1888, 2015.
- [18] A. Bonifati *et al.*, “Designing data marts for data warehouses,” *ACM Trans. Softw. Eng. Meth.*, vol. 10, no. 4, pp. 452–483, 2001.
- [19] R. Kimball and M. Ross, *The data warehouse toolkit: the complete guide to dimensional modeling*. John Wiley & Sons, 2011.
- [20] S. Jupp *et al.*, “A new ontology lookup service at EMBL-EBL,” in *Proc. 8th SWAT4LS Int. Conf.* CEUR-WS.org, 2015, pp. 118–119.
- [21] C. A. Davis *et al.*, “The encyclopedia of DNA elements (ENCODE): data portal update,” *Nucleic Acids Res.*, vol. 46, no. D1, pp. D794–D801, 2017.
- [22] M. Masseroli *et al.*, “Modeling and interoperability of heterogeneous genomic big data for integrative processing and querying,” *Methods*, vol. 111, pp. 3–11, 2016.
- [23] A. Møller, “dk.brics.automaton – finite-state automata and regular expressions for Java,” 2017, <http://www.brics.dk/automaton/>.
- [24] K. Thompson, “Programming techniques: Regular expression search algorithm,” *Commun. ACM*, vol. 11, no. 6, pp. 419–422, 1968.
- [25] M. O. Rabin and D. Scott, “Finite automata and their decision problems,” *IBM J. Res. Dev.*, vol. 3, no. 2, pp. 114–125, 1959.
- [26] S. Ceri *et al.*, “What you always wanted to know about Datalog (and never dared to ask),” *IEEE Trans. Knowl. Data Eng.*, vol. 1, no. 1, pp. 146–166, 1989.
- [27] J. Malone *et al.*, “Ten simple rules for selecting a bio-ontology,” *PLoS Comput. Biol.*, vol. 12, no. 2, p. e1004743, 2016.
- [28] P. L. Whetzel *et al.*, “Bioportal: enhanced functionality via new web services from the national center for biomedical ontology to access and use ontologies in software applications,” *Nucleic Acids Res.*, vol. 39, no. suppl_2, pp. W541–W545, 2011.
- [29] M. Martínez-Romero *et al.*, “NCBO Ontology Recommender 2.0: an enhanced approach for biomedical ontology recommendation,” *J. Biomed. Semant.*, vol. 8, no. 1, p. 21, 2017.
- [30] A. Bernasconi *et al.*, “Ontology-driven metadata enrichment for genomic datasets,” in *Proc. 11th SWAT4LS Int. Conf.* CEUR-WS.org, 2018, pp. 1–10.
- [31] M. Martínez Romero *et al.*, “Using association rule mining and ontologies to generate metadata recommendations from multiple biomedical databases,” *Database*, vol. 2019, 06 2019, baz059.
- [32] F. Albrecht *et al.*, “Deepblue epigenomic data server: programmatic data retrieval and analysis of epigenome region sets,” *Nucleic Acids Res.*, vol. 44, no. W1, pp. W581–W586, 2016.
- [33] L. Nanni *et al.*, “PyGMQL: scalable data extraction and analysis for heterogeneous genomic datasets,” *BMC Bioinf.*, vol. 20, no. 1, p. 560, 2019.
- [34] D. J. Rigden and X. M. Fernández, “The 27th annual nucleic acids res. database issue and molecular biology database collection,” *Nucleic Acids Res.*, vol. 48, no. D1, pp. D1–D8, 2020.
- [35] Gomez-Cabrero *et al.*, “Data integration in the era of omics: current and future challenges,” *BMC Syst. Biol.*, vol. 8, no. Suppl 2, p. I1, 2014.
- [36] V. Lapatas *et al.*, “Data integration in biological research: an overview,” *J. Biol. Res. Thessaloniki*, vol. 22, no. 1, p. 9, 2015.
- [37] D. Smedley *et al.*, “The BioMart community portal: an innovative alternative to large, centralized data repositories,” *Nucleic Acids Res.*, vol. 43, no. W1, pp. 589–598, 2015.
- [38] I. Rodchenkov *et al.*, “Pathway Commons 2019 Update: integration, analysis and exploration of pathway data,” *Nucleic Acids Res.*, vol. 48, no. D1, pp. D489–D497, 2020.
- [39] P. Artimo *et al.*, “ExPASy: SIB bioinformatics resource portal,” *Nucleic Acids Res.*, vol. 40, no. W1, pp. W597–W603, 2012.
- [40] E. L. Hong *et al.*, “Principles of metadata organization at the ENCODE data coordination center,” *Database*, vol. 2016, 03 2016, baw001.
- [41] D. Bujold *et al.*, “The international human epigenome consortium data portal,” *Cell Syst.*, vol. 3, no. 5, pp. 496–499, 2016.
- [42] J. Zhang *et al.*, “The international cancer genome consortium data portal,” *Nat. Biotechnol.*, vol. 37, no. 4, pp. 367–369, 2019.
- [43] M. A. Musen *et al.*, “The center for expanded data annotation and retrieval,” *J. Am. Med. Inform. Assoc.*, vol. 22, no. 6, pp. 1148–1152, 2015.
- [44] A. J. Gray *et al.*, “Bioschemas: From potato salad to protein annotation,” in *Proc. 16th Int. Semant. Web Conf. (Poster Track)*, 2017.
- [45] N. V. Kovalevskaya *et al.*, “DNAdigest and repositiv: connecting the world of genomic data,” *PLoS Biol.*, vol. 14, no. 3, p. e1002418, 2016.
- [46] S.-A. Sansone *et al.*, “DATS, the data tag suite to enable discoverability of datasets,” *Sci. Data*, vol. 4, p. 170059, 2017.
- [47] O. Pastor, “Understanding the human genome: a conceptual modeling-based approach,” in *Proc. 21st Int. Conf. DEXA*. Springer, 2010, pp. 467–469.
- [48] L. Wang *et al.*, “Biostar models of clinical and genomic data for biomedical data warehouse design,” *Int. J. Bioinform. Res. Appl.*, vol. 1, no. 1, pp. 63–80, 2005.
- [49] J. F. R. Román *et al.*, “Applying conceptual modeling to better understand the human genome,” in *Proc. 35th Int. Conf. ER*. Springer, 2016, pp. 404–412.
- [50] A. L. Palacio *et al.*, “A method to identify relevant genome data: Conceptual modeling for the medicine of precision,” in *Proc. 37th Int. Conf. ER*. Springer, 2018, pp. 597–609.
- [51] G. Rambold *et al.*, “Meta-omics data and collection objects (MODCO): a conceptual schema and data model for processing sample data in meta-omics research,” *Database*, vol. 2019, 01 2019, baz002.
- [52] D. Oliveira *et al.*, “Where to search top-k biomedical ontologies?” *Briefings Bioinf.*, vol. 20, no. 4, pp. 1477–1491, 2019.
- [53] C. Jonquet *et al.*, “A system for ontology-based annotation of biomedical data,” in *Proc. Int. Work. DILS*. Springer, 2008, pp. 144–152.
- [54] N. H. Shah *et al.*, “Ontology-driven indexing of public datasets for translational bioinformatics,” *BMC Bioinf.*, vol. 10, no. 2, p. S1, 2009.
- [55] C. B. Giles *et al.*, “ALE: Automated Label Extraction from GEO metadata,” *BMC Bioinf.*, vol. 18, no. 14, p. 509, 2017.
- [56] J. D. Fernandez *et al.*, “Ontology-based search of genomic metadata,” *IEEE/ACM Trans. Comput. Biol. Bioinf.*, vol. 13, no. 2, pp. 233–247, 2016.



Anna Bernasconi earned her Masters in Computer Engineering in 2015 from Politecnico di Milano and University of Illinois at Chicago. Currently PhD candidate at Politecnico di Milano, she works in the field of data-driven genomic computing. Her research interests include bioinformatics data and metadata integration methodologies to support complex biological queries answering.



Arif Canakoglu received his PhD Degree in Computer Engineering in 2016 from Politecnico di Milano and he is currently a Post-Doc fellow at Politecnico di Milano, working in the field of data-driven genomic computing and machine learning applications in the bioinformatics area. His research interests include databases, ontologies, big data processing, cloud computing, machine learning, and bioinformatics.



Marco Masseroli is Associate Professor at the Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB) of Politecnico di Milano, Italy. His research interests are in the area of bioinformatics, focused on distributed Internet technologies, biomolecular databases and ontologies to effectively retrieve, analyze, and integrate genomic information with clinical and genomic data. He is the author of more than 200 publications in international journals, books and conference proceedings.



Stefano Ceri is Professor at the Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB) of Politecnico di Milano. His research has been generally concerned with extending database technology; he has authored over 350 publications. He received two advanced ERC Grants, on Search Computing and on Data-Driven Genomic Computing (GeCo, 2016–2021). He received the ACM-SIGMOD Innovation Award (2013) and is an ACM Fellow.