

Enhancing Clinical Data Warehouses with Provenance and Large File Management: The gitOmmix Approach for Clinical Omics Data

Maxime Wack^{1,2,3,4}, Adrien Coulet^{1,2}, Anita Burgun^{1,2,5}, Bastien Rance^{1,2,3,*}

¹ Centre de Recherche des Cordeliers, Inserm, Université Paris Cité, Sorbonne Université, Paris, France

² Inria Paris, Paris, France

³ Department of Biomedical Informatics, Hôpital Européen Georges Pompidou, AP-HP, Paris, France

⁴ Centre Hospitalier National d'Ophtalmologie des Quinze-Vingts, IHU FOrESIGHT, 75012 Paris, France

⁵ Imagine Institute, Inserm UMR 1163, Université Paris Cité, Paris, France

* corresponding author: `bastien.rance@aphp.fr`

Abstract

Background Clinical data warehouses (CDWs) are essential in the reuse of hospital data in observational studies or predictive modeling. However, state-of-the-art CDW systems present two drawbacks. First, they do not support the management of large data files, what is critical in medical genomics, radiology, digital pathology, and other domains where such files are generated. Second, they do not provide provenance management or means to represent longitudinal relationships between patient events. Indeed, a disease diagnosis and its follow-up rely on multiple analyses. In these cases no relationship between the data (*e.g.*, a large file) and its associated analysis and decision can be documented.

Method We introduce GITOMMIX, an approach that overcomes these limitations, and illustrate its usefulness in the management of medical omics data. GITOMMIX relies on *(i)* a file versioning system: git, *(ii)* an extension that handles large files: git-annex, *(iii)* a provenance knowledge graph: PROV-O, and *(iv)* an alignment between the git versioning information and the provenance knowledge graph.

1 Introduction

Background With the rise of personalized medicine, patient omics data such as RNA or whole genome sequencing (WGS) enrich traditional clinical data, and by consequence find their place in electronic health records (EHR) and clinical data warehouses (CDW) [1, 2, 3]. In this perspective, CDWs enriched with omics data offer an alternative to prospective cohorts for translational studies, *i.e.*, studies typically searching for genotype – phenotype associations such as genetic profiling of sub-groups of diseases or drug responses [4, 5].

CDW-based translational platforms present two main drawbacks that motivated this work. The first drawback is the lack of management of data provenance. These platforms record patient events (such as observations, interventions, decisions) about patients in chronological order, but they seldom explicitly record historical relationships between these events. Accordingly, the questions “what are the observations that supported this decision?”, and inversely “what decisions were made from this observation?” can not be answered by these systems. The second drawback is the lack of management of large data files. Relating a clinical decision, such as a diagnosis, to the content of a large file, such as the files of a whole-genome sequencing, remains difficult with CDWs. However, these two functionalities are crucial for the management of clinical omics studies.

In computer science, *data provenance* is defined as the documentation of where data comes from, and how it was transformed [6]. Among other aspects, provenance facilitates reproducibility in research [7], *i.e.* the ability to obtain the same results by applying the same procedure to the same data [8]. For this reason, standards and tools for data provenance have been developed [9, 10] and widely adopted in fields such as bioinformatics [11, 12], but only parsimoniously diffused to medical informatics. However, provenance and reproducibility are crucial for applications such as clinical decision support tools and their successful transfer to clinical practice. This is particularly true when results are generated from prone-to-error biotechniques, potentially requiring several runs before confirming their validity. One reason for the lack of data provenance management in medical informatics is its absence from CDW systems. Most successful CDW models and their implementation, such as the i2b2 star model [13], the OMOP Common Data Model (CDM) [14], or the eHOP model in France [15] are not supporting detailed provenance.

The *management of large files*, *i.e.*, larger than several hundred megabytes, or longer than 10 thousand lines, is also limited in CDWs. This is mainly due to their use of relational database management systems, which are not designed to handle large files. When such files are supported, they are usually stored aside from the CDW, and the CDW stores an unique reference such as an URL to the file. This is prone to inconsistencies and missing data, as file location and availability rely on independent file management systems, not synchronized with the CDW.

Objective and motivation Our objective is to design an approach that allows storing large data files involved in clinical diagnoses and decisions, the relations between these data and diagnoses and decisions, as well as potential relations between clinical diagnoses and decisions; and to provide ways to query those relations and access the underlying data.

A common issue is to identify links between facts. For example, identifying patients with a liver metastasis within a cohort of lung cancer patients. This group of patients is not simply the set of patients with both diagnoses, as proof of the causal link between the original tumor and metastasis is necessary. Our approach should enable querying that specific relation unambiguously as well as retrieving the supporting data (*e.g.*, the digital pathology images of the primary lesion), and analyses (*e.g.*, the search for variants in sequencing data associated with the disease progres-

sion tracking) for that relation. Accordingly, it would help identifying patients satisfying complex inclusion criteria by querying the CDW in a more clinically meaningful way.

A core requirement of translational research is to access data obtained from high-throughput experiments and associated clinical data. Our approach should enable finding all the observations related to a condition and its longitudinal follow-up, as well as retrieving the corresponding data. More generally, it would allow the query of longitudinal information to access follow-up results or decision (or inversely to past data that motivated a decision).

Proposed solution GITOMMIX allows provenance tracing, large file management, and the encoding of longitudinal relationships in CDWs, by combining: (i) the file versioning system git and its git-annex extension to manage large file histories, (ii) a knowledge graph to encode provenance metadata, (iii) a data model providing an alignment between these two systems, mapping data with metadata.

The rest of the article is organized as follow: Section 2 presents the building bricks of our approach ; Section 3 presents GITOMMIX itself ; Section 4 illustrates its use for the management of clinical omics studies.

2 Material

2.1 Semantic Web tools for data and provenance

The Semantic Web proposes a set of standards and tools that facilitate sharing, linking, and processing data by associating them with formally defined semantics [16]. This work relies on three Semantic Web standards: RDF (Resource Description Framework) [17], SPARQL (SPARQL Protocol And RDF Query Language), and PROV-O (PROV Ontology). RDF, the Semantic Web standard for encoding knowledge graph, is a data model that represents data in the form ⟨subject, predicate, object⟩ triples, to describe a binary relation associating a subject and an object. SPARQL is a query language for RDF knowledge graphs [18]. PROV-O is a standard ontology recommended since 2013 by the W3C for the encoding of provenance metadata [10]. PROV-O is built around three main concepts: **Entities**, **Activities**, and **Agents**. Entities represent physical or virtual objects, such as data sets or atomic elements of data. Entities can be generated or modified by Activities. Activities are realized by Agents. Entities can also be directly attributed to Agents. (Figure 1a)

Adopting Semantic Web technologies provides additional tools contributing to the adherence to the FAIR principles [19].

2.2 git and git-annex

git is a distributed open-source file versioning system created in 2005 to support the Linux kernel development and now ubiquitously used in software development [20]. git traces historical changes within files in a directory, called a *repository*. It uses a directed acyclic graph (DAG) structure, the *git graph*, to record repository states, called *commits*. Because repositories are distributed and thus need to follow independent changes in various locations, branching and merging of histories is permitted, and is a core mechanism of collaborative development in software engineering. File additions, removals, or modifications are recorded in commits, which are accompanied by a *commit message* describing the changes assigned to a *commit author*. Commits are uniquely identified in

a repository by a cryptographic hash code, which can be seen as a signature of the content of the commit. Any commit in the history of a repository and its associated files can be retrieved from the corresponding unique commit hash code.

git has originally been designed to trace changes in source code files, usually relatively small text files, but does not scale to large files. git-annex, a third-party extension to git, has been created to overcome this limitation and handle large files [21]. git-annex stores the designated files contents aside from the git repository and takes over the management of those files, while still recording the historical information by tracing a reference to the file within git. It provides its own operations for adding and retrieving files, supporting a range of popular efficient file-hosting back-ends. The stored reference is a cryptographic hash of the file content, making git-annex a *content-addressable* file store: any change in a file has the consequence of modifying its cryptographic hash, enabling the unique identification of multiple versions of the same file.

3 Methods - git O m m i x

We designed GITOMMIX with three main components:

- a *data model* that records and semantically links clinical data and decisions that are related in term of provenance,
- a *system* that traces changes in data, pointing at their up-to-date clinical interpretations,
- an *association* between the data model and the system to ensure a progressive encoding of data provenance, at the time of data changes.

We defined a set of operations to build, manage, and query patient data history represented with the GITOMMIX data model.

The GITOMMIX *data model* uses the PROV-O concepts of **Entities**, **Agents**, and **Activities**, and various possible relations between those concepts.

We use them to represent clinical data and their provenance relations: the **Agent** concept represents data providers, which can either be a human or a machine; the **Activity** concept represents analyses, software runs or other methods that produce one data element; the **Entity** concept represents any data element recorded in a CDW, associated or not with files. We extend entities into five subtypes: patients, samples, data, results, and diagnoses. Patients and samples are considered as data elements because in the context of a CDW, they are indeed identifier of patients or samples. Sample is a general naming encompassing identifiers of biological samples, but also of images or audio recordings. Diagnoses can be any kind of clinical decision, but we restrict our study to diagnoses only, for simplicity.

The most central relation of PROV-O, linking **Entities** together, is **wasDerivedFrom**. The relation **wasRevisionOf** is also used in the specific case of derivations that are data modifications. Figure 1a illustrates these concepts and relations, and their use to represent data elements of a CDW and how one might derive from another. This derivation is a *many-to-many* relation, as a sample can generate multiple data elements, and multiple results can lead to a single diagnosis. This relation between two entities is the atomic block that is repeated to build sequences of data elements derived from a patient, as illustrated in Figure 1b. In our model, a diagnosis **wasDerivedFrom** a result, which **wasDerivedFrom** some data, which **wasDerivedFrom** a sample, which **wasDerivedFrom** a patient. In this figure and throughout the rest of the article, we adopt the PROV-O prescribed

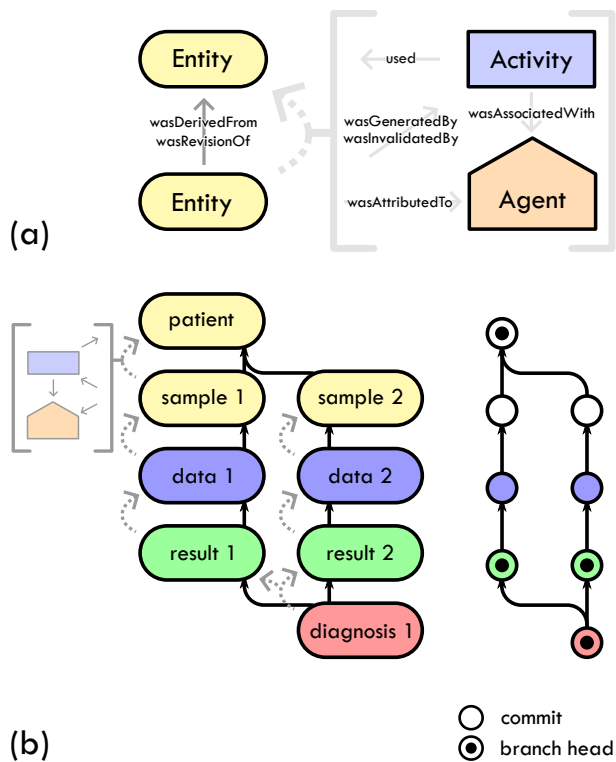


Figure 1: The GITOMMIX data model. (a) Derivation (or revision) relationship between two data elements of a CDW, represented with concepts and relationships of the PROV-O ontology. (b) Sequence of data elements deriving one from each other, as a repeat of the pattern showed in (a). The bottom left schema illustrates the correspondence between a sequence represented with PROV-O (on the left) and a git commit graph (on the right).

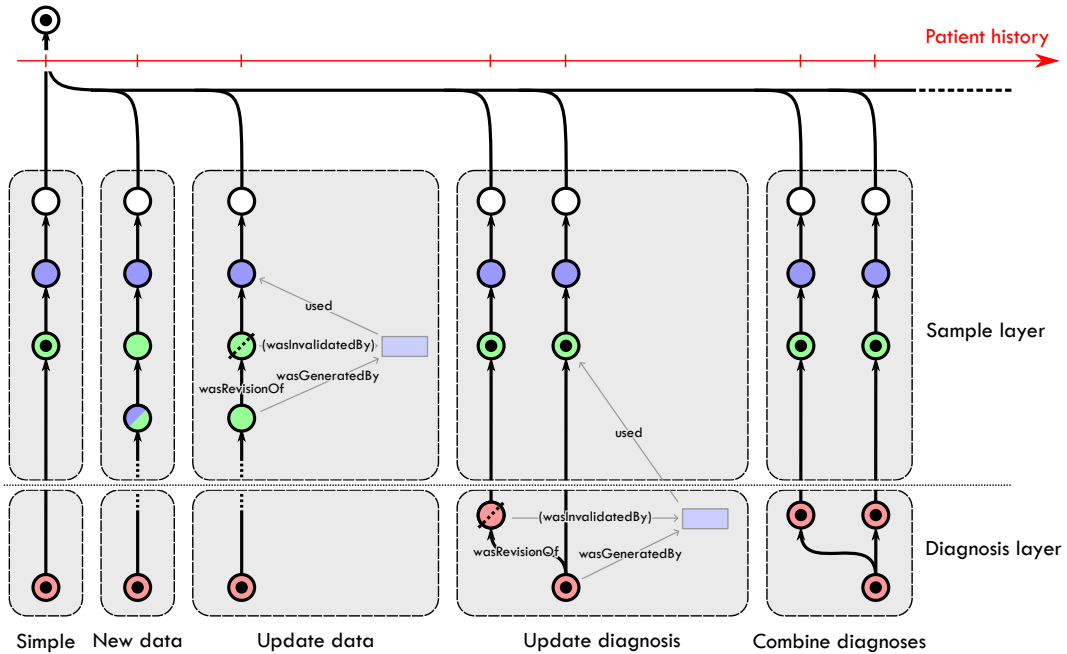


Figure 2: Overview of the possible events composing a patient history in GITOMMIX. Nodes are commits, edges commit branching associated with a provenance type from PROV-O. Default type of relation is `wasDerivedFrom`, except when otherwise specified. Blue, green and red are data, results and diagnoses, respectively.

shapes to distinguish between Entities, Activities and Agents. In addition we use different colors to distinguish entities: blue for data, green for results, and red for diagnoses. Other entities are kept blank.

For a more concrete example, a diagnosis of diabetes (an ICD10 code in the CDW), was derived from a laboratory result of high blood glucose concentration (a LOINC code), which in turn was derived from a blood glucose analysis (identified by an internal lab number), which was derived from a blood sample (a nursing procedure code).

In the specific case of modifications, updates or invalidations of diagnosis, results or data, the `wasDerivedFrom` relation is replaced by `wasRevisionOf`. Providers and methods can optionally be added to further document the derivation relationship between entities. Accordingly, an entity E_1 `wasAttributedTo` a provider P , which `wasAssociatedWith` a method M . E_1 `wasGeneratedBy` M , which `used` a previous entity E_2 . For example, E_1 is a WGS assay attributed to a lab technician P , themselves associated with a short read sequencing ‘, which generated the sequence files E_1 using the patient sample E_2 .

Implementing patient data history with git In GITOMMIX, we propose to rely on the git versioning system to trace patient data elements and associated clinical decisions. Each patient is represented with its own single git repository, the git commit graph progressively built with new data, changes, and interpretations. Small and large data files associated with patients observa-

tions are referenced and relationships between data elements and decisions are implemented using commits and git branching mechanisms. To facilitate maintaining this structure, we describe two “layers”: a *sample* layer for the histories of data derived from every sample acquired from the patient (*e.g.*, a biological sample, an image, or an audio file), and a *diagnosis* layer for the relations between clinical decisions. These two layers are illustrated in Figure 2.

The sample layer encompasses *sample branches*. Each new sample acquisition is materialized by a new git branch in the patient’s git graph. Each data and result derived from a sample is sequentially added to that sample branch as new commits. Multiple revisions of these data or results can be added to a sample branch, possibly invalidating a previous version.

The second layer encompasses *diagnosis branches*. Clinical diagnoses are materialized by new branches on top of the sample branches, following a different construction rule. Diagnosis commits can derive from one or multiple results (thus from multiple sample branches) using the git branch merging function. Such a merge represents the joint contribution of multiple results to a single diagnosis. Diagnoses can be further revised, combined, or invalidated by new merges of new results or diagnoses.

In git, branches are pointers to the latest commit in that branch, called the *HEAD*. In a sample branch, the HEAD points to the most up-to-date information and data related to that sample. In a diagnosis branch, the HEAD points to the most up-to-date diagnosis.

Alignment between the data model and patient git graph Sequences of provenance relations represented with the data model can be aligned to git commit graphs, as illustrated in Figure 1b. In this alignment, each PROV-O Entity corresponds to a commit in the git graph, and derivation relationships between Entities corresponds to parent-child commit relationships.

We implement this alignment by reusing the structure git offers to associate metadata to commits. Indeed, every commit has an author, a date, and a message composed of a subject and a body. We use the author and date to record the provider and date, respectively. The message subject records the entity type and its id in the following form: `type:id` (*e.g.*, `patient:123467890`, `diagnosis:ICD10_I10`); and the message body records the associated metadata, encoded in turtle RDF. For example, when adding a biopsy sample to a patient, the following RDF pattern is written into the commit message body:

```
:sample:{sampleId}
  a prov:Entity ;
  a :sample ;
  prov:generatedAtTime xsd:dateTime:{date} ;
  prov:wasDerivedFrom :patient:{patientId} ;
  rdf:label "{sampleId}" ;
  prov:wasAttributedTo :provider:{providerId} ;
  prov:wasGeneratedBy :method:biopsy .

:provider:{providerId}
  a prov:Agent ;
  a :provider .

:method:biopsy
  a prov:Activity ;
```



```

a :method ;
prov:startedAtTime xsd:dateTime:{date} ;
prov:wasAssociatedWith :provider:{providerId} ;
prov:used :patient:{patientId} .

```

Using this commit metadata, the formal representation of provenance is preserved and closely associated with the corresponding data files, with the relationships between entities mirrored in the git graph structure. Concatenating all the commit message bodies of a particular git history builds the corresponding RDF knowledge graph by incrementally adding nodes and relations. The resulting knowledge graph has the advantage of offering query and reasoning facilities beyond those provided by git alone.

Examples of patient data histories Figure 2 illustrates the possible events in a patient history and their representation in GITOMMIX .

The ‘Simple’ box in Figure 2 illustrates the trivial case of a diagnosis obtained from a single sample and a single biomedical analysis. A sample branch is created (blank node); a data file is added with a commit (blue node); a result is added with a new commit (green node). Next, a diagnosis is added on top of the result by the creation of a new diagnosis branch (red node). Target nodes are the HEAD of their respective branch.

The ‘New data’ box in Figure 2 illustrates the case of new data or results obtained from the same sample. Those are added sequentially, accumulating the information produced from a single sample in the same history. The graph structures of the RDF and git graphs can slightly differ here, as the git history stays linear while the RDF graph splits, as data always derives from the sample.

The ‘Update data’ box in Figure 2 illustrates the case of data or results updating or replacing previous ones. Those are added sequentially, as in the previous case, with the use of the `wasRevisionOf` relation instead of `wasDerivedFrom`. In the case of invalidation, the invalidated entity is additionally documented with a temporal relationship `invalidatedAtTime`, and optionally with a relationship to the method it `wasInvalidatedBy`. Note that multiple entities of the same kind can be invalidated at once by a single new entity. In all those cases, accessing the repository “at” a diagnosis gives access to all the files in the version that was used to lead to that diagnosis.

The ‘Update diagnosis’ and ‘Combine diagnoses’ boxes in Figure 2 illustrate when a diagnosis emerges from the combination of multiple exam results, possibly originating from multiple samples, and from previous partial or erroneous diagnoses. This is achieved in git using the merge operation to combine entities contributing to this diagnosis. In both the update and combine cases a new diagnosis branch is created. The update diagnosis case is useful for aggregating more information related to a single diagnosis, or combining previous partial or symptomatic diagnoses into an etiological diagnosis. The combined diagnoses case can be used to signal that multiple co-existing diagnoses are related through a syndromic diagnosis. The new diagnosis can either stay unchanged when new analysis stay compatible with that diagnosis, or a different one derived from the additional information.

Figure S1 in Supplementary Data shows the actual git graph (a) and PROV graph (b) of the implementation of these operations in GITOMMIX . A script listing all the GITOMMIX commands needed to build this repository is provided with the software package.




Figure 3: The three types of queries enabled by GITOMMIX: (a) Provenance, (b) Up-to-date, and (c) Clinical timeline queries.

4 Results

We implemented GITOMMIX as a set of operations and queries that can be called from a command line interface. These commands are mostly abstractions over the underlying git commands managing the repositories.

Editing git graphs and associated provenance GITOMMIX allows users to formulate simple commands such as “add the NGS assay files of this patient’s lung biopsy”. Each command triggers a series of git operations and enriches the commit metadata with provenance. Note that GITOMMIX only enriches the git graph and metadata and does not suppress any of it, following the philosophy of version control systems. Accordingly, GITOMMIX provides operations to add, revise, or invalidate elements (patients, samples, data, results, diagnoses). These operations can be combined to enable more complex ones, such as “add the variant calling result derived from the NGS assay of this lung biopsy, and makes it revise the inconclusive pathology report produced earlier”.

Queries with gitOmmix GITOMMIX provides three main types of queries, illustrated in Figure 3. (1) retrieving the provenance of any entity in the patient history. This is supported by the simple fact that visiting a repository “at” a commit accesses all the files accumulated up to that commit. Large files stored in the annex are downloaded on demand only. git logging facilities enable provenance to be further narrowed to specific entities, time periods, providers, etc. Figure 3(a) illustrates this type of query listing all the data that contributed to an input diagnosis.

(2) retrieving the most up-to-date data, results, or diagnoses of a branch. As each sample and diagnosis is represented as a branch containing its whole history, navigating to its *HEAD* returns the most recent version of data, results, and diagnoses, as illustrated Figure 3(b). Chained with the first operation, this enables retrieving the additional data that documents the evolution of a diagnosis, *e.g.*, a pathology report that documents the recurrence of a lung cancer.

(3) returning the timeline of a patient’s successive diagnoses. Because the RDF representation of the provenance is represented in a piece-wise manner in the metadata of each entity, any subgraph of provenance can be built by concatenating a selection of pieces. The resulting RDF can in turn be returned, queried in SPARQL, or used to produce graph visualizations. The timeline operation

displayed in Figure 3(c) is implemented with a SPARQL query that only selects diagnoses and diagnosis-diagnosis relations.

Indeed, in addition to these three common queries, GITOMMIX supports running arbitrary SPARQL queries on patients histories.

A real-world clinical case We consider the previously published case report of a patient with a metastatic HPV-induced high grade anal intraepithelial neoplasia (HGAIN) [22, 23]. We schematized the git ommix graph associated with this patient history in Figure 4: (1) The patient was diagnosed with an HGAIN during high-resolution anoscopy that led to the realization of a biopsy and pathologist confirmation of the diagnosis, completed with PCR identification of HPV6, 11, and 16. (2) The HGAIN was subsequently surgically resected, with confirmed free resection margins. (3) Two years later, the patient presented feverish back pains accompanied by weight loss, that were misdiagnosed at first. After a couple of months and a new feverish episode, a new bone biopsy conducted the pathologist to conclude to a diagnosis of squamous cell carcinoma (SCC) metastasis. Further investigations using immunohistochemical and PCR assays detected HPV16 DNA in the biopsy tissue. This allowed linking the vertebral lesion to the HGAIN despite the absence of other signs of anal SCC contemporary to the metastasis diagnosis. (4) Furthermore, as the patient had been participating in a research protocol that included the collection of plasma samples, those samples were retrospectively analyzed using quantitative digital droplet PCR, showing the presence of HPV16 circulating DNA in increasing blood concentration between the two diagnoses. (5) Further investigations using HPV capture and NGS showed that the exact same HPV16 subvariant was detected throughout all samples, from the initial HGAIN to the vertebral metastasis.

The actual PROV graph generated for this example is shown in Figure S2 of Supplementary Data.

Implementation An implementation of GITOMMIX is available at <https://www.github.com/gitOmmix/gitOmmix> as Open Source Software. It is implemented in bash and offers a user interface in the form of a command-line tool which includes help and auto-completion. It relies on the command-line versions of git, git-annex, the rapper and roquet command-line tools from <https://librdf.org> for RDF file management, and graphviz [24] for graph visualization.

5 Discussion

Integration to clinical data warehouses GITOMMIX is designed to enrich CDWs to enable the support of data provenance and large files. Entities within a patient’s repository are uniquely identified by an automatically generated identifier (a SHA1 hash of the git commit introducing the entity). Thus, the pair (patient id, unique hash) is an unambiguous reference to an object into an instance of GITOMMIX. Accordingly, GITOMMIX entities can be referenced in a CDW by associating the corresponding unique hashes to observations (*e.g.*, the sample associated to the surgical procedure, the data file associated with a biological analysis, the full result report associated with an image exam, the diagnosis associated with a stay in the billing system). In observation-based models such as i2b2, this can be achieved by re-using an existing column to store the associated hash. In other scenarios where no free column is available, adding a specific column to the schema for this purpose is sufficient and does not interfere with the CDW. In CDWs using the OMOP CDM, although a similar mechanism could be used, a cleaner implementation

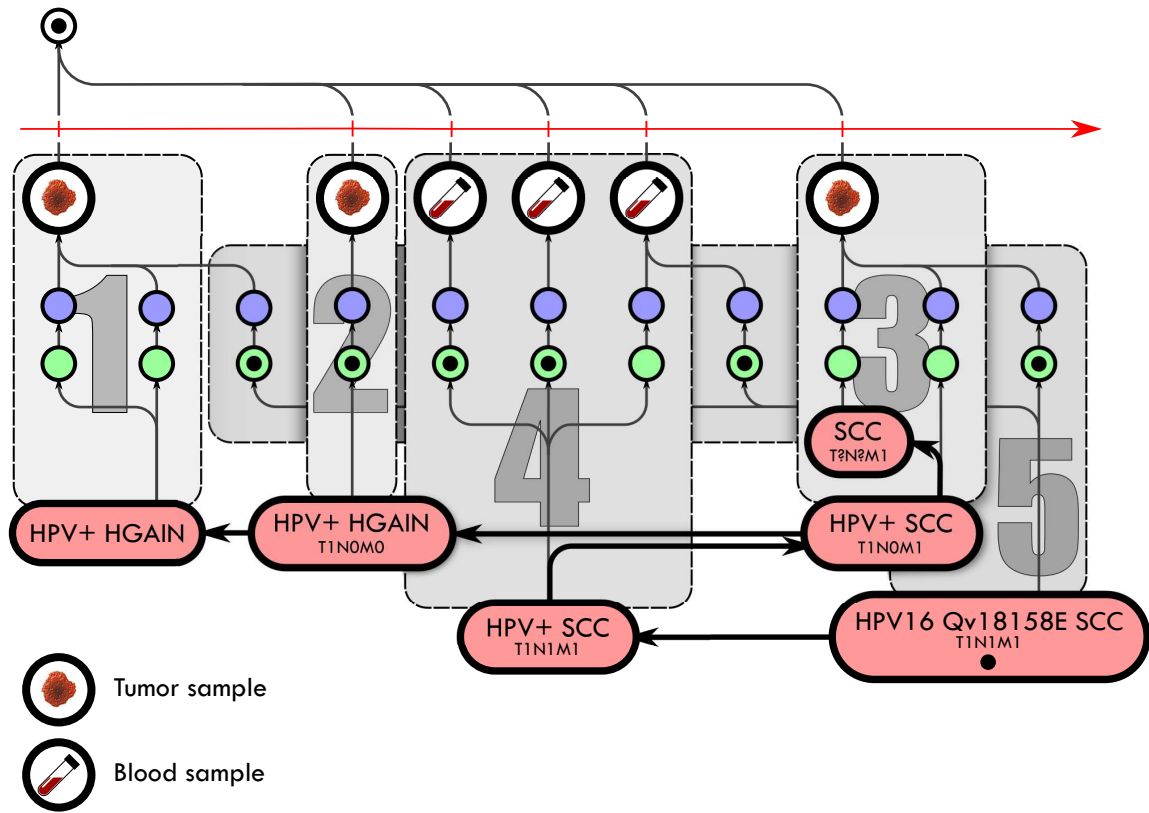


Figure 4: Example of a patient's history in GITOMMIX: (1) The patient is diagnosed with an HGAIN from the examination of a biopsy, with evidence of HPV6, 11, and 16 infection. (2) The HGAIN is resected with free resection margins and as such is coded as T1N0M0. (3) Later, a vertebral lesion is identified as a SCC metastasis of unspecified origin then linked to the HGAIN after evidence of HPV16 in the vertebral lesion. (4) HPV16 ctDNA in increasing quantities is retrospectively detected on biobanked plasma samples. (5) A specific variant of HPV16 in its episomal form is evidenced throughout all samples

using a new concept (*e.g.*, “gitommix_hash”) and a `fact_relationship` linking the observation to its `GITOMMIX` hash would be preferable. It is consequently possible to navigate back and forth between `GITOMMIX` and the CDW as necessary, filtering on patient identifier and commit hash in the observations table on the CDW side or targeting the commit hash on the `GITOMMIX` side.

Related works In a 2017 article [25], Murphy *et al.* described three methods for combining clinical and genomic data within the i2b2 CDW. The first one involves integrating genomic results as structured data using the Sequence Ontology. The second one uses the i2b2/transSMART platform with its *ad hoc* ontology and data storage mechanism. The third one uses a NoSQL database containing functionally annotated results, linked to i2b2 *via* a custom i2b2 module. All those methods involve transformations of the genetic results, making it impossible to access the primary data; and do not embed links between results, making it impossible to track longitudinal relationships between assays. Our solution is architecturally similar to the third method, in that it adds an external layer to the CDW. However, it is more independent as it does not rely on a specific CDW implementation and does not necessitate in-depth adaptation of CDW, but only relies on common tools.

`GITOMMIX` is compliant with any particular data schema or controlled vocabulary. It could be added seamlessly to the first described method using the Sequence Ontology, linking each structured result to its source data.

Various initiatives have tentatively added structured representations of genomic data in the OMOP CDM, such as the genomic CDM (G-CDM) [26]. This complements `GITOMMIX` by enabling further structured representation of data hosted in `GITOMMIX`.

Limitations The current implementation of `GITOMMIX` is a proof of concept and for this reason presents some limitation. First, it is local and centralized, and does not yet support the management of shared and remote repositories, as permitted by `git` and `git-annex` and intended for `GITOMMIX`. Second, its interface is still rudimentary and limited to experienced users.

Advantages However, our solution has multiple advantages over previously described systems, by giving the ability to integrate arbitrary large data and by enabling the representation of relations between these data points. It does so without relying on an entirely new paradigm around data representation in CDW, or needing heavy adaptations to the system currently in use. It acts as a plug-in solution, agnostic from the CDW system in use, and is supported by standards and tools that were originally designed to address the specific issues at hand: keeping record of file histories, managing large files, and tracing provenance in a formal way.

`GITOMMIX`, as well as the tools it relies on, is open source backed by open standards. This allows to readily benefit from the capabilities of these tools, enabling powerful file management, and clearly defined semantics, reasoning capabilities, and interoperability. Although our system prescribes a general structure to its data model, it does not restrict the usage of additional features from the underlying systems. For example, the base `PROV-O` triples generated for each entity can be supplemented at will to construct richer provenance graphs. `git-annex` supports a diverse collection of file storage back-ends to host arbitrary large files, locally or remotely, on-premise or in the cloud, offering to benefit from efficient storage solutions.

Regarding scalability, as `git-annex` separates the management of large file from the management of the repository, repositories themselves stay very light in term of memory and thus responsive

to query. And because each patient exists as its own git repository, operations can inherently be parallelized by running as many GITOMMIX processes as needed.

Perspectives By enriching data with provenance metadata and enabling access to versioned source data files, GITOMMIX contributes to a better adherence to FAIR principles in the management of complex clinical data. In particular, it ensures findability by assigning persistent and unambiguous identifiers, providing rich metadata, and a standard way to search within this metadata; accessibility by making source data available through git; interoperability by relying on standard knowledge representations; and reusability by adding detailed provenance and allowing access to all versions of the data.

For these reasons GITOMMIX allows reproducibility and consistency in conducting translational studies, particularly retrospective studies based on large data that are more and more routinely collected during care. In addition, it may also benefit clinical care as it documents clinical decisions explicitly and in a FAIR format [3].

On the technical side, using established standards and tools allows for the addition of features supported by those tools. For example, git enables authors to cryptographically sign their commits, which could be used to add a layer of security to the tracing of provenance. git repositories can contain references to other git repositories using *submodules*. Using submodules could allow even richer provenance tracing by directly referencing the actual analysis code, pipeline, or tool at the version in which it was used to produce observations.

6 Conclusion

We introduce GITOMMIX, a relatively simple and lightweight system combining semantic web, file versioning, and content-addressable distributed file storage to represent and manage provenance and large source data in clinical data warehouses. It includes all the functions required to build a patient's history graph and store associated files, navigate and query history using SPARQL, and retrieve the specific files related to any event. We base our proposition on widely accepted systems and a model leveraging the shared DAG structure underlying these systems.

We provide a proof of concept implementation demonstrating feasibility and practical use of GITOMMIX, and illustrate its use with real-world use case about diagnosis based on clinical omics data. It is open to contributions and will be extended to support additional functions.

Abbreviations

CDW: Clinical Data Warehouse
ctDNA: circulating DNA
DAG: Directed Acyclic Graph
ddPCR: digital droplet Polymerase Chain Reaction
EHR: Electronic Health Record
FAIR: Findable, Accessible, Interoperable, Reusable
G-CDM: Genomic Common Data Model
HGAIN: High Grade Anal Intraepithelial Neoplasm
HPV: Human Papilloma Virus
i2b2: Informatics for Integrating Biology and the Bedside

ICD-10: International Classification of Diseases, 10th revision
LOINC: Logical Observation Identifiers Names and Codes
NLP: Natural Language Processing
NoSQL: non-SQL
OMOP CDM: Observational Medical Outcomes Partnership Common Data Model
OWL: Web Ontology Language
POC: Proof of Concept
PROV-O: Provenance Ontology
RDF: Resource Description Framework
RDFS: RDF Schema
SCC: Squamous Cell Carcinoma
SHA-1: Secure Hash Algorithm 1
SPARQL: SPARQL Protocol and RDF Query Language
WGS: Whole Genome Sequencing

Authors contributions

MW: conceptualization, software, visualization, writing (original draft, review and editing).
AC: supervision, conceptualization, writing (review and editing).
AB: writing (review and editing), validation, funding acquisition.
BR: supervision, conceptualization, funding acquisition, writing (editing and review)

Funding resources

We benefit from a government grant managed by the Agence Nationale de la Recherche under the France 2030 program, reference ANR-22-PESN-0007 ShareFAIR.

Acknowledgments

Dr. H el ene P er e and Dr. David Veyer for the fruitful discussions about their inspiring clinical projects. Linus Torvalds for creating linux and git, on which this contribution is based.

References

- [1] V. Canuel, B. Rance, P. Avillach, P. Degoulet, A. Burgun, Translational research platforms integrating clinical and omics data: a review of publicly available solutions, *Briefings in Bioinformatics* 16 (2) (2015) 280–290.
- [2] A. Trifan, J. L. Oliveira, Patient data discovery platforms as enablers of biomedical and translational research: A systematic review, *Journal of Biomedical Informatics* 93 (2019) 103154. doi:10.1016/j.jbi.2019.103154.
- [3] A. J. Robertson, A. J. Mallett, Z. Stark, C. Sullivan, It Is in Our DNA: Bringing Electronic Health Records and Genomic Data Together for Precision Medicine, *JMIR Bioinformatics and Biotechnology* 5 (1) (Jun. 2024). doi:10.2196/55632.

- [4] S. Pendergrass, M. D. Ritchie, Phenome-wide association studies: leveraging comprehensive phenotypic and genotypic data for discovery, *Current Genetic Medicine Reports* 3 (2) (2015) 92–100.
- [5] O. Lyudovyk, Y. Shen, N. P. Tatonetti, S. J. Hsiao, M. M. Mansukhani, C. Weng, Pathway analysis of genomic pathology tests for prognostic cancer subtyping, *Journal of Biomedical Informatics* 98 (2019) 103286. doi:<https://doi.org/10.1016/j.jbi.2019.103286>.
- [6] R. Bose, J. Frew, Lineage retrieval for scientific data processing: a survey, *ACM Comput. Surv.* 37 (2005) 1–28.
- [7] S. N. Goodman, D. Fanelli, J. P. A. Ioannidis, What does research reproducibility mean?, *Science Translational Medicine* 8 (341) (2016) 341ps12–341ps12. arXiv:<https://www.science.org/doi/pdf/10.1126/scitranslmed.aaf5027>, doi:10.1126/scitranslmed.aaf5027.
- [8] F. C. Y. Benureau, N. P. Rougier, Re-run, Repeat, Reproduce, Reuse, Replicate: Transforming Code into Scientific Contributions, *Frontiers in Neuroinformatics* 11 (2017) 69. doi:10.3389/fninf.2017.00069.
- [9] R. Ikeda, A. D. Sarma, J. Widom, Logical provenance in data-oriented workflows?, in: 2013 IEEE 29th International Conference on Data Engineering (ICDE), IEEE, 2013, pp. 877–888.
- [10] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, PROV-O: The PROV Ontology, *W3C recommendation* 30 (2013).
- [11] R. Almeida, W. da Silva, K. Castro, M. E. Walter, A. Araujo, M. Holanda, S. Lifschitz, AProvBio: An architecture for data provenance in bioinformatics workflows using graph database, in: 2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), 2017, pp. 2139–2144. doi:10.1109/BIBM.2017.8217989.
- [12] V. W. Salazar, J. V. F. Cavalcante, D. de Oliveira, F. Thompson, M. Mattoso, BioProv - A provenance library for bioinformatics workflows, *Journal of Open Source Software* 6 (67) (2021) 3622. doi:10.21105/joss.03622.
- [13] S. N. Murphy, G. Weber, M. Mendis, V. Gainer, H. C. Chueh, S. Churchill, I. Kohane, Serving the enterprise and beyond with informatics for integrating biology and the bedside (i2b2), *Journal of the American Medical Informatics Association : JAMIA* 17 (2) (2010) 124–130. doi:10.1136/jamia.2009.000893.
- [14] E. A. Voss, R. Makadia, A. Matcho, Q. Ma, C. Knoll, M. Schuemie, F. J. DeFalco, A. Londhe, V. Zhu, P. B. Ryan, Feasibility and utility of applications of the common data model to multiple, disparate observational health databases, *J. Am. Med. Inform. Assoc.* 22 (3) (2015) 553–564.
- [15] J. Madec, G. Bouzillé, C. Riou, P. Van Hille, C. Merour, M.-L. Artigny, D. Delamarre, V. Raimbert, P. Lemordant, M. Cuggia, ehop clinical data warehouse: From a prototype to the creation of an inter-regional clinical data centers network., *Studies in Health Technology and Informatics* 264 (2019) 1536–1537.

- [16] T. Berners-Lee, J. Hendler, O. Lassila, The semantic web, *Scientific American* 284 (5) (2001) 34–43.
- [17] W3C, Resource Description Framework (RDF) Model and Syntax Specification, <https://www.w3.org/TR/PR-rdf-syntax/Overview.html>, visited Aug. 22, 2024.
- [18] T. W. S. W. Group, Sparql 1.1 overview, <https://www.w3.org/TR/sparql11-overview/>, visited Aug. 9, 2023 (2013).
- [19] 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, *Scientific data* 3 (1) (2016) 1–9.
- [20] L. Torvalds, Initial revision of “git”, the information manager from hell, README file, visited Aug. 9, 2023 (2005).
- [21] git-annex web page:, <https://git-annex.branchable.com/>, visited Aug. 9, 2023.
- [22] H. Péré, R. Vernet, S. Pernot, J. Pavie, N. Robillard, J. Puech, S. Lameiras, M.-L. Lucas, A. Nicolas, C. Badoual, B. Rance, L. Bélec, L. Weiss, M. Wack, D. Veyer, Episomal HPV16 responsible for aggressive and deadly metastatic anal squamous cell carcinoma evidenced in peripheral blood, *Scientific Reports* 11 (1) (2021) 4633. doi:10.1038/s41598-021-84110-2.
- [23] D. Veyer, J. Pavie, S. Pernot, M. Mandavit, S. Garrigou, M.-L. Lucas, L. Gibault, V. Taly, L. Weiss, H. Péré, HPV-circulating tumoural DNA by droplet-based digital polymerase chain reaction, a new molecular tool for early detection of HPV metastatic anal cancer? A case report, *European Journal of Cancer* (Oxford, England: 1990) 112 (2019) 34–37. doi:10.1016/j.ejca.2019.02.012.
- [24] E. R. Gansner, S. C. North, An open graph visualization system and its applications to software engineering, *Software: Practice and Experience* 30 (11) (2000) 1203–1233. doi:[https://doi.org/10.1002/1097-024X\(200009\)30:11<1203::AID-SPE338>3.0.CO;2-N](https://doi.org/10.1002/1097-024X(200009)30:11<1203::AID-SPE338>3.0.CO;2-N).
- [25] S. N. Murphy, P. Avillach, R. Bellazzi, L. Phillips, M. Gabetta, A. Eran, M. T. McDuffie, I. S. Kohane, Combining clinical and genomics queries using i2b2 - Three methods, *PloS One* 12 (4) (2017) e0172187. doi:10.1371/journal.pone.0172187.
- [26] S. J. Shin, S. C. You, Y. R. Park, J. Roh, J.-H. Kim, S. Haam, C. G. Reich, C. Blacketer, D.-S. Son, S. Oh, R. W. Park, Genomic Common Data Model for Seamless Interoperation of Biomedical Data in Clinical Practice: Retrospective Study, *Journal of Medical Internet Research* 21 (3) (2019) e13249. doi:10.2196/13249.