

Extracting, Transforming and Archiving Scientific Data

Daniel Lemire¹ and Andre Vellino²

¹ Université du Québec à Montréal,
100 Sherbrooke West
Montreal, QC, H2X 3P2 Canada
lemire@gmail.com

² CISTI - National Research Council of Canada
Building M-55, 1200 Montreal Road
Ottawa, ON K1A 0S2 Canada
andre.vellino@nrc.ca

Abstract. It is becoming common to archive research datasets that are not only large but also numerous. In addition, their corresponding metadata and the software required to analyse or display them need to be archived. Yet the manual curation of research data can be difficult and expensive, particularly in very large digital repositories, hence the importance of models and tools for automating digital curation tasks. The automation of these tasks faces three major challenges: (1) research data and data sources are highly heterogeneous, (2) future research needs are difficult to anticipate, (3) data is hard to index. To address these problems, we propose the Extract, Transform and Archive (ETA) model for managing and mechanizing the curation of research data. Specifically, we propose a scalable strategy for addressing the research-data problem, ranging from the extraction of legacy data to its long-term storage. We review some existing solutions and propose novel avenues of research.

1 Introduction

The conventional workflow model of science—whereby the scientist proposes a hypothesis, devises a series of experiments, performs the experiments, generates data and produces a publication in a peer-reviewed journal—is no longer adequate to characterize present-day scientific endeavours. First, a significant amount of scientific research is devoted to experimental design, data-collecting, developing increasingly precise measurement techniques and managing the acquired data. Furthermore, researchers today have an increasing ability to share resources and methods and a greater need to handle large volumes of data. They also have more opportunity to collaborate across a variety of disciplines and have a greater diversity of channels for disseminating results, data and software beyond conventional publication channels.

Hence there is a need for tools that automate the curation processes beyond merely storing and archiving large volumes of research data. They also need to enable data reuse, interoperability and discovery. This challenge is especially difficult because research communities differ so widely in their needs and practices [23] that universally applicable conventions are impossible to establish.

Furthermore, to create complete data archives we must also be able to extract data from previously published “backfiles” whose legacy data content may not have ever been managed, curated or archived at all, let alone with discovery, reuse and repurposing in mind.

To achieve these objectives, we present a new data-management model for digital libraries which addresses the problems of large scale automation of extraction, transformation and archiving (ETA) of scientific research data. Our proposal is founded on a mature model — Extract Transform Load (ETL) — that has been developed for business data warehousing [12] and complements the data management elements of existing digital curation models such as the Digital Curation Centre (DCC) Lifecycle Model [11]. Because we favor automatization when possible, our approach is founded on *sheer curation* [15]: the curation activities are integrated within the normal workflow of those creating the data.

We review the ETL model in §2 and show how it can be adapted to the problem of extracting data in §3; transforming it in §4 and archiving it in §5.

2 ETL

ETL is a process model used in data warehousing to integrate heterogeneous data sources and enable uniform data analytics. The “Extract” component of the ETL process aims at harvesting data from disparate sources in a variety of formats. The “Transform” part of the ETL process performs cleaning operations and applies encoding rules to convert the source data into a more coherent form. The “Load” phase takes the transformed data that conforms to a uniform data schema and makes it available to a database system on which, for example, analysis tools can be executed.

In the Enterprise Database marketplace, software tools, such as Oracle Warehouse Builder, DB2 Warehouse Edition and Microsoft SQL Server Integration Services implement and automate the ETL model. These tools pay particular attention to the enterprise needs of performance and scalability as well as the requirements for data migration and auditing.

The parallels with the requirements for managing scientific research data on a large scale are clear: research data, even within the same scientific discipline manifests in a variety of heterogeneous formats and there is a present need for “data harmonization”. The scale and distribution of research-data also means that the models and methods used in pure data warehousing should apply.

3 Data Extraction

In their study of the life cycle of e-Science data, Wallis et al. [5] identified the following phases: (1) Experimental Design, (2) Calibration, (3) Data Capture or Generation, (4) Data cleaning and Derivation, (5) Data Integration, (6) Data Derivation, (7) Data Analysis, (8) Publication, Storage and Preservation. They found that scientists need to access their data at each phase and must be able to use and integrate data from multiple sources. As the authors point out

The lack of an integrated framework for managing these types of scientific data presents significant barriers not only to those scientists conducting the research, but also to those who would subsequently reuse the data.

To enable reuse, we must publish the data. Until the Web became ubiquitous, the data itself was rarely published separately from the research articles, and for a good reason: data cannot be understood without a context. Yet decoupling the data from its context is invaluable because it enables verification, reuse and re-purposing.

There are several strategies to decouple the data while retaining its link to its context. The most prevalent approach is to require researchers to upload their data to a curated repository after publication of the corresponding research articles. For example, DiLauro et al. [9] describe a system wherein the data is captured during the submission of the research article. This ensures that the data is properly linked to the research article and that data submission is part of the researchers' workflow.

A preferable alternative is to systematically archive the data as it is being collected and processed [7], or even as it is being acquired by instruments such as some data repositories in astronomy do when the data is collected by telescopes. Finally, the data can be extracted from the research articles or reports themselves: from the tables of results, from the results section—commonly found in the abstract in medical articles—or elsewhere in the document.

Besides the data itself, we must also capture metadata to help users retrieve, assess and reuse the data. Decoupled data may then be linked to a region of text—such as the text which surrounds a table in a research article. Such text serves as indexable metadata, as HTML text does for images on web indexes.

3.1 Extraction from Legacy Sources

Useful data may be inconveniently embedded in a variety of previously published documents. For example, until recently, it was common to store data on paper as plots. Thus, researchers are now forced to recover data by scanning plots from research papers [19].

Research data is also published as tables in PDF, HTML or XML documents. Thankfully, automated data extraction systems such as Tableseer and the SciVerse Applications platform allow researchers to search for and extract tables embedded within documents. Other publishers—such as the Public Library of Science (PLOS)—make available the content of all their journals in XML, a machine readable format that makes it convenient to extract data from published articles using an XQuery engine. Hence, massive numbers of documents can be processed automatically with little effort.

3.2 Decoupled and Linked Data

Decoupled data needs to meet only three criteria:

- It must be free from the confines of the publication format of a research article (PDF, DOC, HTML). That is, it should be in a data-appropriate format that enables further machine processing (CSV, XML or SQL assertions).
- It must be reusable. Thus, it must be available, complete, licensed for reuse and documented. If appropriate, it should follow known data models and schemas [13].

- It must be possible to refer to it to specifically and independently of any research article. For example, it could have a unique identifier.

There is a distinction but also a relationship between the concepts of “linked data” and “decoupled data”. Linked data exposes, shares, and connects pieces of data using URIs and RDF [4]. Hence decoupled data may become linked data. Indeed, decoupling the data from its textual source makes subsequent linking possible.

4 Transformation

Scientists usually transform their data before archiving it:

Mergers and joins Researchers routinely integrate data sets from different sources to derive indicators and measures: astronomers may combine the data from several telescopes and geophysicists may combine satellite data with ground sensors. A frequent, but mostly implicit, type of join occurs when mostly static data is used as part of a derivation. For example, physical constants or geographical data is often used in conjunction with recently collected data.

Data cleaning Almost all research data requires cleaning. The collected data might be inconsistent or contradictory. Outliers indicating faulty measures are common. A particular challenge in science is baseline correction. For example, climatologists need to correct the temperature records for the effect of growing cities.

Data filtering and aggregation It is common for scientists to record more data points than needed: this extra data must be either filtered out or aggregated. When medical researchers process electrocardiograms (ECGs), they routinely keep as little as only the location of one data point per heart beat (e.g., the location the R wave). Geophysicists may carry aeromagnetic survey using planes that record several samples per second, whereas they are ultimately only interested in a geological map having a relatively low resolution.

Data mapping A common mapping in science is a change of units (e.g., from inches to cm). Numerical data can be rounded (e.g., to 3 significant digits).

This list is by no means exhaustive. Other transformations include compression, deduplication and validation [1]. Moreover, scientists increasingly work with data sets so large that they cannot manually inspect them. We must rely on algorithms. Thankfully, there are user-friendly tools to help users transform their data more reliably [8]. In this respect, we find Google Refine [8] particularly interesting.

4.1 Formatting and Standardization

An important type of transformation is the one that maps the data between different formats. For example, long term archival may require a machine-independent format such as netCDF [20] whereas, for on-line access, it might be preferable to have the data in an SQL format.

Beyond the data itself, the metadata must also be properly formatted for interoperability and long-term storage. For example, the Core Scientific Metadata Model (CSMD) [18] is generic enough to apply to a variety of disciplines but also detailed enough to enable the reuse and repurposing of data within and across scientific disciplines.

5 Archiving

We distinguish three types of data which may require archival:

- Raw data, which might result directly from an experiment or a simulation, or it might have been extracted from legacy sources.
- Derived data, which is the result of any processing on the raw data, including cleaning (correction for errors). It includes data integration wherein various data sets are used to create a new data set.
- Resultant data, which is the final product, typically what might be published by the authors along with their research article.

After several decades of manual curation, scientific data repositories such as GenBank [3] offer a wealth of raw data and associated metadata, including references to the published and gray literature. There are even journals such as *Earth System Science Data* dedicated to publishing raw data.

We know from this experience that a proper data archive must support data embargoes [7] and must provide access control. This is especially of concern if researchers upload their data prior to the final publication of their research articles. Sometimes the data needs to be archived and accessible while remaining partially confidential.

Moreover, a data archive should support versioning: even within a single team, there might be several versions of the same data set [24]. For long-term storage, data must be protected against loss and corruption as well as malicious attacks [2]. Data sets should be properly indexed and documented and should have unique identifiers.

While it might be tempting to only store the resultant data, there are at least two problems with this approach:

- other researchers may mistrust transformations that they cannot verify;
- it is difficult to predict how and in what format the data might be most useful to others, even with the best intentions.

Thus, publishing only the resultant data may limit its reusability [17]. Moreover, as Yan et al. [25] report, “scientists are highly motivated to publish the entire data trail along the analysis pipeline.”

Data archiving with an ETA framework affords an opportunity to do more than mere curation. The association of research data with other artefacts such as research articles makes it possible to automate the analysis of metadata to discover trends in the published literature. Thus it should be possible to measure whether progress toward knowledge objectives have been achieved. Similarly, such metadata analysis could detect anomalies and inconsistencies among research results. Citation data analysis could be the basis for recommending research data sets and research papers [21] to researchers. Last but not least it is also possible to mine the data itself to discover novel results [16].

5.1 Specialization or Integration?

In conventional data warehousing, experts distinguish between data marts, which are specialized domain-specific data repositories (e.g., for accounting) and integrated data warehouses, which provide a uniform layer of abstraction from the data-domain. In the context of scientific data repositories, the Australian National Data Service (ANDS) is an example of the latter whereas GenBank is an example of the former.

While domain-specific repositories such as GenBank are simpler to setup and maintain, they are not as conducive to interdisciplinary research. Even though they are more difficult to implement, integrated data warehouses afford a greater likelihood for the data to be repurposed across different disciplines. They are more likely to persist over time and are easier for machines to resolve outside of the subject domain context, such as the database system that generated them. An example of this difference is found in the data-identification mechanism. GenBank assigns its own unique dataset identifier (accession numbers) based on subject-domain conventions for referencing genes and proteins. However, for interdisciplinary research, it might be preferable to refer to those same datasets with a domain-independent persistent identifier system such as DOIs for data sets [22] granted by institutions affiliated with DataCITE [6]. One advantage of adopting a domain-independent identifier is that it eliminates the need for multiple methods for name-resolution.

The trade-off between the two methodologies (centralized versus specialized) is well documented in the data warehousing literature [14]. Experience suggests that the integration of data marts may lead to metadata inconsistencies while the integrated data warehouse approach is costly and difficult to initiate. Specialized data marts tend to be more concise as only the information deemed relevant by the community is included. Integrated data warehouses are more likely to rely on homogeneous technology: they use fewer software vendors. Data marts are often more dynamic: it is comparatively easier to add new feature or new metadata when you must only address the needs of a specific community.

5.2 A Diverse Software Architecture

Many bibliographic repositories—of either text or data—suffer from a common ailment that could be addressed within our proposed ETA process. Indeed, data archiving is often performed by relational databases whose core concepts were invented in the 1960s, and whose technology is insufficient with respect to modern information retrieval needs: semantic search, question answering, content-clustering and dynamic schemas, to name only a few. For these purposes, full-text indexing tools such as Apache Lucene that can do full text analysis e.g., stemming, part-of-speech tagging, term-frequency analysis are starting to replace databases. Similarly, document-oriented databases such as CouchDB and MongoDB might offer the necessary flexibility to dynamically support many different database schemas corresponding to different domains, with relatively little maintenance.

6 Conclusion and Future Research

The management of scientific data repositories can benefit from lessons learned in data warehousing [12]. Both specialized data marts and integrated data warehouses have a role to play in the data archiving ecosystem. When integrating heterogeneous data is too difficult, the data mart approach is preferable. Otherwise, An integrated data warehouse approach favours interdisciplinary collaboration: it offers uniform metadata conventions, persistent identification nomenclature and better automation for ingest. However, if the sources of heterogeneous data are too diverse, the domain-specific data mart approach may be preferable. Moreover, the choice might be guided by the available funding: in most instances, the integrated approach will prove more expensive and require more time.

In either case, data cannot be routinely processed in an ad hoc fashion. The ETA process must be automated as much as possible. While some data repositories can develop their own automation architecture, there is an opportunity to develop more generic ETA tools. For example, we could extend existing open source ETL tools such as Talend³ or Pentaho Data Integration (PDI)⁴.

Donoho et al. [10] recommend making available both the data and the instructions necessary to reproduce any published figure or other published objects. While journals and funding agencies may require that the data is available, as yet we lack conventions on how to document and archive the transformation from raw data to, for example, a figure. Ideally, the data used to generate a figure or a table should always be available through a permanent identifier such as a Data DOI.

Extracting, transforming and archiving heterogeneous data can accommodate a diversity of software architectures. Expertise in data warehousing require the ability to integrate a wide variety of technologies, data formats and data models. We cannot expect to index all research data using only a few simple models: data archives must embrace diversity. To cope with this diversity, we need extensible data-management tools.

An essential insight is that managing the flow of data, from its extraction to its storage and retrieval, is often more important than merely curating the provided data [25]: the life-span of raw data may also include derived and resultant data. Furthermore, it is necessary to consider that data may originate from extraction processes and appropriately transformed and identified.

References

1. M. Altman. A fingerprint method for scientific data verification. In T. Sobh, editor, *Advances in Computer and Information Sciences and Engineering*, pages 311–316. Springer Netherlands, 2008.
2. G. Antunes, J. Barateiro, M. Cabral, J. Borbinha, and R. Rodrigues. Preserving digital data in heterogeneous environments. In *JCDL '09*, pages 345–348, 2009.
3. D. Benson, M. Boguski, D. Lipman, J. Ostell, and B. Ouellette. GenBank. *Nucleic acids research*, 26(1):1, 1998.

³ <http://www.talend.com/>

⁴ http://www.pentaho.com/products/data_integration/

4. C. Bizer, T. Heath, and T. Berners-Lee. Linked data-the story so far. *International Journal on Semantic Web and Information Systems*, 5(3):1–22, 2009.
5. C. L. Borgman, J. C. Wallis, M. S. Mayernik, and A. Pepe. Drowning in data: digital library architecture to support scientific use of embedded sensor networks. In *JCDL '07*, pages 269–277, 2007.
6. J. Brase. Datacite - a global registration agency for research data. *International Conference on Cooperation and Promotion of Information Resources in Science and Technology*, 0:257–261, 2009.
7. M. Cragin, C. Palmer, J. Carlson, and M. Witt. Data sharing, small science and institutional repositories. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1926):4023, 2010.
8. R. Cyganiak, F. Maali, and V. Peristeras. Self-service linked government data with dcat and gridworks. In *Proceedings of the 6th International Conference on Semantic Systems*, I-SEMANTICS '10, pages 37:1–37:3, New York, NY, USA, 2010. ACM.
9. T. DiLauro, M. Cyzyk, E. Metsger, and M. Patton. Capturing and curating published data. In *JCDL '10*, pages 399–400, 2010.
10. D. Donoho, A. Maleki, I. Rahman, M. Shahram, and V. Stodden. Reproducible research in computational harmonic analysis. *Computing in Science & Engineering*, pages 8–18, 2009.
11. S. Higgins. The DCC Curation Lifecycle Model. *The International Journal of Digital Curation*, 3:134–140, 2008.
12. W. H. Inmon. The data warehouse and data mining. *Commun. ACM*, 39:49–50, November 1996.
13. G. Janée, J. Mathena, and J. Frew. A data model and architecture for long-term preservation. In *JCDL '08*, pages 134–144, 2008.
14. N. Jukic. Modeling strategies and alternatives for data warehousing projects. *Commun. ACM*, 49:83–88, April 2006.
15. S. Macdonald and L. Martinez-Urbe. Collaboration to data curation: Harnessing institutional expertise. *New Review of Academic Librarianship*, 16(1):4–16, 2010.
16. D. MacMillan. NextBio. *Reference Reviews*, 23(2):39–40, 2009.
17. T. Malik, L. Nistor, and A. Gehani. Tracking and Sketching Distributed Data Provenance. In *IEEE International Conference on e-Science '10*, 2010.
18. B. Matthews, S. Sufi, D. Flannery, L. Lerusse, T. Griffin, M. Gleaves, and K. Kleese. Using a Core Scientific Metadata Model in Large-Scale Facilities. *International Journal of Digital Curation*, 5(1), 2010.
19. M. Mayersohn and S. Tannenbaum. On reclaiming data from the literature: literature data “R and R” (recovery and reanalysis). *American Journal of Pharmaceutical Education*, 62:363–370, 1999.
20. R. Rew and G. Davis. NetCDF: an interface for scientific data access. *IEEE Computer Graphics and Applications*, 10(4):76–82, 2002.
21. K. Sugiyama and M.-Y. Kan. Scholarly paper recommendation via user’s recent research interests. In *JCDL '10*, pages 29–38, 2010.
22. A. Treloar and R. Wilkinson. Access to Data for eResearch: Designing the Australian National Data Service Discovery Services. *International Journal of Digital Curation*, 3(2), 2008.
23. J. C. Wallis, M. S. Mayernik, C. L. Borgman, and A. Pepe. Digital libraries for scientific data discovery and reuse: from vision to practical reality. In *JCDL '10*, pages 333–340, 2010.
24. J. C. Wallis, A. Pepe, M. S. Mayernik, and C. L. Borgman. An exploration of the life cycle of escience collaborative data. In *iConference '08*, 2008.
25. E. Yang, B. Matthews, and M. Wilson. Enhancing the Core Scientific Metadata Model to Incorporate Derived Data. In *IEEE International Conference on e-Science '10*, 2010.