



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Curating Covid-19 Data in Links

Citation for published version:

Galpin, V & Cheney, J 2021, Curating Covid-19 Data in Links. in *Provenance and Annotation of Data and Processes*. Lecture Notes in Computer Science, vol. 12839, Springer, pp. 237-243, 8th International Provenance and Annotation Workshop, 19/07/21. https://doi.org/10.1007/978-3-030-80960-7_19

Digital Object Identifier (DOI):

[10.1007/978-3-030-80960-7_19](https://doi.org/10.1007/978-3-030-80960-7_19)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Provenance and Annotation of Data and Processes

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Curating Covid-19 data in Links

Vashti Galpin¹[0000–0001–8914–1122] and James Cheney^{1,2}[0000–0002–1307–9286]

¹ University of Edinburgh, UK
{Vashti.Galpin, James.Cheney}@ed.ac.uk
² The Alan Turing Institute, UK

Abstract. Curated scientific databases play an important role in the scientific endeavour and support is needed for the significant effort that goes into their creation and maintenance. This demonstration and case study illustrate how curation support has been developed in the Links cross-tier programming language, a functional, strongly typed language with language-integrated query and support for temporal databases. The chosen case study uses weekly released Covid-19 fatality figures from the Scottish government which exhibit updates to previously released data. This data allows the capture and query of update provenance in our prototype. This demonstration will highlight the potential for language-integrated support for curation to simplify and streamline prototyping of web-applications in support of scientific databases.

1 Introduction

Curated scientific databases take significant human effort to develop and then to maintain [2]. These databases are crucial in providing knowledge bases for science to proceed and users of the data need to be able to trust in their contents. One way to support this trust is for users to be able to review provenance information.

We consider update provenance [2] as a specific use case, based on the fact that regularly released datasets can include updates to previously released datasets, and our case study is based on data that has been released weekly by the Scottish government on Covid-19 deaths [8].

We are developing a prototype curation interface using Links³, a cross-tier programming language with language-integrated query that is being extended with experimental temporal database features [3].

Both the prototype and Links’s temporal features are works-in-progress, and as we develop the interface, we are considering the more general issues that would apply to any curated scientific databases where update provenance is to be captured using temporal databases.

2 Links

Links is a cross-tier programming language developed at the University of Edinburgh [3]. It is *cross-tier* in the sense that it allows a developer to write a single type-checked

³ <https://links-lang.org/>

program which that can then be executed efficiently on multi-tier Web architectures. Typically, a web application consists of three distinct applications coded in different languages: HTML and JavaScript for the browser, a different language for the application logic and SQL for the database queries, leading to complex coordinating issues. Links provides a single language in which to implement these three aspects, thus simplifying the development process. For our curation interface, the built-in translation to JavaScript completely abstracts away the details and there is no JavaScript code in the source Links program.

Furthermore, Links supports *language-integrated query* whereby database queries can be written in Links itself resulting in safe and efficient queries. In particular when writing database-related code, the support provided by type inference for row types in Links makes referring to the fields of interest in a table while abstracting from the others straightforward. This leads to compact and readable programming style. Links requires the types of the database tables to be declared explicitly in the program, which enables checking of language-integrated queries and supports the temporal query transformations. This adds some work at the beginning and if there is schema evolution, work is needed re-synchronize these declarations. Support for schema evolution in Links is a future area of research. Links currently supports PostgreSQL, MySQL 5.x and SQLite 3. Links covers a wide range of SQL but some aspects are ongoing research such as support for set and bag semantics [9].

In terms of use for curated scientific databases, IUPHAR/BPS Guide to PHARMACOLOGY (GtoPdb) [1] has been reimplemented in Links as a case study to demonstrate that Links is suitable for this task both in terms of functional correctness and performance [4]. However, because of its size and complexity, GtoPdb is not a good candidate for exploring the strengths and weaknesses of Links's new temporal support; instead in this demonstration we present the current Covid-19 curation prototype, which is the next step in developing curation functionality for this type of database in Links. We consider both the interface as the end product, and how the development of the interface is supporting by the features of Links.

3 Temporal databases

Temporal databases provide the ability to record when a row in a table is valid, either with respect to the database itself using a transaction time period, or with respect to the real world using a valid time period [5]. We can use this functionality to track the validity of the data items that are updated. Figure 1 illustrates how an update differs between a standard database on the left, and a temporal database on the right. In the former, the data is replaced and the previous values are lost. However, in the temporal case, additional columns to record the start and end of the time period of validity form part of the key and allow previous values to be recorded. The interpretation of the validity is dependent on the particular application.

Although temporal extensions to SQL have been standardized [6], many current popular relational database implementations have no built-in support for temporality, although it is possible to implement temporal tables by the use of explicit additional fields and complex hand-generated SQL queries [10]. In comparison, Links now pro-




Fig. 1. Standard database update (left) compared with temporal database update (right)

vides support for transaction-time tables and valid-time tables, allowing both sequenced (taking time period into account) and non-sequenced queries over these tables. This is achieved by interacting with a standard relational database such as PostgreSQL, using the approach currently provided by Links for generating standard queries and translating temporal queries into standard queries as described above [10], thereby avoided the need for the application developer to generate these complex SQL statements. An example is given in the next section.

4 Curation functionality

To investigate how to develop curation functionality in Links, we choose a dataset that raised questions of interest around updating of data. To set the context for discussion of the type of queries that can be considered, we first provide more information about the data we used in implementing the prototype.

4.1 The case study

We identified data that was released weekly by the Scottish Government through the National Records of Scotland website [8], as suitable and interesting for the prototype exploration. Each week, since early in the pandemic a spreadsheet (as a CSV file and Excel document) has been released (amongst other data) with counts of Covid-19 fatalities for individual weeks and data categories including sex, age, health board, local authority, and location. An example of a data category is “Location of death” and its subcategories are “Care Home”, “Home/Non-institution”, “Hospital” and “Other institution”. Our application supports upload of this CSV data and its transformation into a temporal table as illustrated in Figure 2.

4.2 Provenance-based queries

One of the aims of this research is to identify the queries that are interesting for update provenance. An obvious query is how an individual data item has changed over time, but other potential queries include finding out whether one category of data has changed more often than another. Dependency is also relevant: are the changes of a data item associated or correlated with changes in other data.




Fig. 2. Uploading new data

Examples of queries that the interface can support (or will support in the future):

1. What was the number of female fatalities in the week of 20 April 2020 when that data was first uploaded
2. What is the current number of female fatalities in the week of 20 April 2020?
3. What is the range of female fatalities in the week of 20 April 2020 over all uploads?
4. Which updates to female fatalities were rejected and when?
5. Have there been more updates to female fatalities or male fatalities considering all weeks?
6. Which health board has seen most updates in the first six months of the pandemic?
7. How are updates for the Lothian health board figures correlated with those for the local authority of Edinburgh?

4.3 The prototype curation interface

The interface has been developed using the Bootstrap HTML, CSS and JavaScript library. We present two screenshots illustrating the functionality of the prototype interface in Figure 3. The left-hand figure shows updates for a specific week that have arisen in a subsequent week. By grouping them together, a user is able to assess the consistency of this update⁴. These updates can be accepted or rejected together, or each can be considered individually in the context of other updates to that data item.

The following shows the Links code which updates the main table of fatalities from a table of accepted updates for counts relating to the Lothian health board.

```
for (x <-- AcceptedUpdates)
  [update sequenced (y <-v- CovidDeaths)
   between (valid_from_date, forever())
   where (x.week==y.week && x.category=="Lothian")
   set (count = x.new_value, file_id=x.new_id)];
```

⁴ The dataset has the feature that it contains counts for categories that subsume other categories. We decided to capture all categories rather than just the minimal ones to support checking for inconsistencies.




Fig. 3. Screenshots: decision making (left) and provenance (right)

The double arrows, \leftarrow and \leftarrow , indicate iteration over the rows of the named tables, with \leftarrow indicated that the table is a temporal table. The keyword `sequenced` indicates that this update should result in the modification of one record and the insertion of a new record as illustrated in Figure 1. To write this query in SQL for a non-temporal database would be more complex, requiring explicit updates and insertions.

The right-hand screenshot in Figure 3 shows a provenance query on a single data item. For the updates on this data item, it is possible to see the updates in the context of other updates that occurred in the same week. Other provenance queries can be done on data categories (such as queries 5, 6 and 7 above), and a similar approach can be applied to weeks. The final menu item allows for queries such as query 4.

The prototype has taken around 120 hours to develop to date, approximately 1 person month of effort by the first author, who was not previously familiar with Links. It consists of around 1800 lines of Links code, as well as supporting CSS and JavaScript code. Reuse of existing code will make the implementation of additional queries less time-consuming.

5 Conclusion

We have developed a prototype curation interface in Links to demonstrate and investigate curation functionality for a selected dataset with interesting features relating to data update. Links, as a cross-tier programming language, has provided a typed and functional language for implementing all of the appearance of the web interface (using the Bootstrap library), the logic of the interface and querying of the database.

By using Links, we have avoided the need to code temporal features in SQL for databases without temporal features, to transition to a database such as MariaDB [7] which has temporality, or to deal with the integration errors that occur during conventional web development using SQL.

Another option would be to use a temporal middleware/stratum translator that maps temporal SQL queries to plain ones [11]. This requires the same kind of information

about the database schema that Links requires, but does not provide the advantages of language integration and type checking that we get from Links.

We have successfully created an interface that supports a number of queries and further work involves determining how to support the generation of a curation interface with an arbitrary database schema where selected tables have temporal features.

Acknowledgements This work was supported by ERC Consolidator Grant Skye (grant number 682315) and by an ISCF Metrology Fellowship grant provided by the UK government's Department for Business, Energy and Industrial Strategy (BEIS). We thank Simon Fowler for his assistance in using the temporal extensions to Links.

References

1. Armstrong, J., Faccenda, E., Harding, S., Pawson, A., Southan, C., Sharman, J., Campo, B., Cavanagh, D., Alexander, S., Davenport, A., Spedding, M., Davies, J.: The IUPHAR/BPS Guide to PHARMACOLOGY in 2020: extending immunopharmacology content and introducing the IUPHAR/MMV Guide to MALARIA PHARMACOLOGY. *Nucleic Acids Research* **48**, D1006–D1021 (2020)
2. Buneman, P., Cheney, J., W.-C.Tan, Vansummeren, S.: Curated databases. In: Proceedings of the Twenty-Seventh ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '08. p. 1–12 (2008)
3. Cooper, E., Lindley, S., Wadler, P., Yallop, J.: Links: Web programming without tiers. In: 5th International Symposium on Formal Methods for Components and Objects, FMCO 2006. pp. 266–296. LNCS 4709 (2006)
4. Fowler, S., Harding, S., Sharman, J., Cheney, J.: Cross-tier web programming for curated databases: A case study. *International Journal of Digital Curation* **16** (2020)
5. Jensen, C., Snodgrass, R.: Temporal database. In: *Encyclopedia of Database Systems*. Springer (2009), https://doi.org/10.1007/978-0-387-39940-9_395
6. Kulkarni, K., Michels, J.: Temporal features in SQL: 2011. *SIGMOD Record* **41**(3), 34–43 (2012)
7. MariaDB: What's new in MariaDB Server 10.3 (2018), <https://mariadb.com/resources/blog/whats-new-in-mariadb-server-10-3/>, last accessed: 23 April 2021
8. National Records of Scotland: Deaths involving coronavirus (COVID-19) in Scotland: archive (2021), <https://www.nrscotland.gov.uk/statistics-and-data/statistics/statistics-by-theme/vital-events/general-publications/weekly-and-monthly-data-on-births-and-deaths/deaths-involving-coronavirus-covid-19-in-scotland/archive>, last accessed: 23 April 2021
9. Ricciotti, W., Cheney, J.: Query lifting - language-integrated query for heterogeneous nested collections. In: 30th European Symposium on Programming, ESOP 2021. pp. 579–606. LNCS 12648 (2021)
10. Snodgrass, R.: *Developing Time-Oriented Database Applications in SQL*. Morgan Kaufmann Publishers Inc. (1999)
11. Torp, K.: Temporal strata. In: *Encyclopedia of Database Systems*. Springer (2009), https://doi.org/10.1007/978-0-387-39940-9_409