## Application Notes

# ReviewR: a light-weight and extensible tool for manual review of clinical records

David A. Mayer[1], Luke V. Rasmussen [2], Christopher D. Roark[3], Michael G. Kahn[4], Lisa M. Schilling[5], and Laura K. Wiley [1]

[1]Department of Biomedical Informatics, Colorado Center for Personalized Medicine, University of Colorado Anschutz Medical Campus, Aurora, Colorado, USA, [2]Department of Preventive Medicine, Northwestern University Feinberg School of Medicine, Chicago, Illinois, USA, [3]Department of Neurosurgery, University of Colorado Anschutz Medical Campus, Aurora, Colorado, USA, [4]Department of Biomedical Informatics, University of Colorado Anschutz Medical Campus, Aurora, CO, USA, and [5]Division of General Internal Medicine, Department of Medicine, Data Science to Patient Value Program, University of Colorado Anschutz Medical Campus, Aurora, Colorado, USA

Corresponding Author: Laura K. Wiley, PhD, Department of Biomedical Informatics, Colorado Center for Personalized Medicine, University of Colorado Anschutz Medical Campus, 1890 N. Revere Court, Aurora, CO 80045, USA; laura.wiley@cuanschutz.edu

### ABSTRACT

**Objectives:** Manual record review is a crucial step for electronic health record (EHR)-based research, but it has poor workflows and is error prone. We sought to build a tool that provides a unified environment for data review and chart abstraction data entry.

**Materials and Methods:** ReviewR is an open-source R Shiny application that can be deployed on a single machine or made available to multiple users. It supports multiple data models and database systems, and integrates with the REDCap API for storing abstraction results.

**Results:** We describe 2 real-world uses and extensions of ReviewR. Since its release in April 2021 as a package on CRAN it has been downloaded 2204 times.

**Discussion and Conclusion:** ReviewR provides an easily accessible review interface for clinical data warehouses. Its modular, extensible, and open source nature afford future expansion by other researchers.

Key words: electronic health records, data warehousing, software, observational studies

### LAY SUMMARY

When doing research using data from electronic health records (EHRs), data may need to be extracted by hand, either to perform the study or to ensure its accuracy. However, many researchers cannot access the EHR for this purpose. Even when researchers have access, they must flip between their review list, the EHR, and the location they are recording the results of their review, which is difficult and can cause errors. We developed a software application, ReviewR, to make this process easier and less error prone and have used it in 2 real-world projects.

# INTRODUCTION

Electronic health records (EHRs) are commonly used for clinical, translational, and population health research.[1–3] Although significant advances have been made in methods for automated data extraction, chart/record review (ie, extracting or abstracting information from a patient's EHR manually) remains a crucial tool for EHR-based research.[4–6] However manual chart review is complex, time consuming, and error prone.[7] Typically, chart review occurs in the source EHR (eg, Epic, Cerner) with the extraction results stored in a secondary system (eg, Excel, REDCap[8]). This leaves researchers juggling multiple application windows which has poor usability and increases the potential for data entry errors (like mistyping the patient ID or entering data on the wrong row of the spreadsheet). Further, many EHRs lack sophisticated search capabilities (eg, using regular expressions to search for multiple relevant terms simultaneously), increasing the complexity of record review. Providing researchers with a single, integrated, and optimized tool would help improve the efficiency and quality of manual record review.

Although traditionally chart review has relied upon using the source EHR, some institutions restrict EHR access to only clinically credentialed staff. Others may require researchers to access deidentified data warehouses or restrict access to only those patients covered under research agreements. This prevents researchers from performing traditional chart review because researchers are restricted to using a data warehouse or specific data extracts. Thus there is a need for a tool to support record review that uses data from clinical data warehouses instead of the traditional EHR interface. A number of tools have been developed to support visualization and searching of clinical data warehouses including i2b2,[9] OHDSI's Atlas,[10] Leaf,[11] and EMERSE (for clinical notes only)[12]; however, these tools do not support integrated chart abstraction. Additionally, these tools all require server-based deployment which typically requires support from data warehouse and/or IT teams. Tools exist for text annotation including the Clinical Language Annotation, Modeling, and Processing Toolkit (CLAMP),[13] Knowtator,[14] and the brat rapid annotation tool,[15] however, these tools are designed to work with individual notes or sentences rather than entire patient records. Given that manual record review often relies on a combination of structured and unstructured data from the EHR, there is an unmet need for a lightweight and flexible tool that has a unified review interface.

# OBJECTIVES

Our primary objective was to design and build a tool that provides a streamlined workflow for performing manual chart review from clinical data warehouses. Specifically, the tool should have: (1) search and filter capabilities that can be applied across an entire patient record, and (2) a unified interface that minimizes context switching by supporting simultaneous review of data and data entry into a chart abstraction form. While these objectives can streamline chart review for all users (including clinical users with access to the source EHR), we specifically wanted to support research and informatics teams who only have access to clinical data warehouses for manual record review. In the development of the tool, we also followed 3 key design principles: (1) adhere to electronic data capture best practices, (2) support lightweight and portable deployment options that allow use by individual researchers or organizations, (3) provide easy extensibility to support different clinical data models and relational database management systems (RDBMS) used to store and query clinical data.

# MATERIALS AND METHODS

ReviewR is an R Shiny[16] application built using modules and the {golem} framework[17] for production grade app development.[18] This approach allows ReviewR to be distributed as a regular R package supporting R v3.5.0 or later.[19] ReviewR was developed and documented using {devtools},[20] the {tidyverse},[21] and a number of other utilities.[22–27] The user interface leverages multiple dashboarding and widget toolkits[28–31] and CSS/JavaScript tools.[32,33] Clinical records are presented using the {DT} package,[34] an R interface to the DataTables JavaScript library,[35] which allows users to easily filter columns and perform complex searches using regular expressions. Support for multiple RDBMSs is provided through {dbplyr}[36]—a package that converts regular dplyr-based code into SQL using a {DBI}[37] mediated database connection. ReviewR currently supports clinical data stored in SQLite,[38] PostgreSQL,[39] and Google BigQuery,[40,41] but can be extended to include any RDBMS supported by {dbplyr}.

Data models are supported in ReviewR using a schema definition (ie, list of table and associated column names) and a matched set of display functions. The data model is automatically detected by selecting the schema with the highest number of table/column name matches. We provide development functions for extending ReviewR to support any custom data model and offer out of the box support for multiple versions of the OMOP common data model[42] and the MIMIC3 database.[43] Chart abstractions are captured in REDCap[8] with ReviewR translating REDCap instruments into native Shiny widgets and transmitting data using the REDCap API and associated R interfaces.[44,45] ReviewR supports multiple instruments and a number of commonly used field types and data validation rules (see Table 1).

To test the extensibility and face validity of ReviewR, we conducted 2 demonstration projects. First, we sought to extend ReviewR to connect to a new RDBMS (Microsoft SQL Server) to support phenotype-based chart review. Second, we sought to extend ReviewR to support a custom data model and compare chart review conducted using ReviewR to chart review performed in the source EHR. One of the authors (CDR) reviewed 50 records in both Epic and ReviewR using a crossover design (ie, 25 records were reviewed first in Epic and then ReviewR, with the other 25 reviewed in Reviewer first) to extract history of intracranial aneurysm, subarachnoid hemorrhage (SAH), and aneurysmal SAH. We then analyzed review concordance.

# RESULTS

ReviewR can be used locally by individual users, or it can support multiple concurrent users when deployed to a Shiny Server. Docker images for both local and server installations may be built from the Dockerfiles included in the package source. ReviewR is available for

**Table 1.** Supported REDCap field and data validation types

| Field type | Validation options |
| --- | --- |
| Checkboxes (multiple answers) | NA |
| Multiple choice—drop-down list (single answer) | NA |
| Multiple choice—radio buttons (single answer) | NA |
| Notes box (paragraph text) | NA |
| Text box (short text, number, date/time, ...) | Date (M–D–Y); Integer |
| True—False | NA |
| Yes—No | NA |

download on CRAN (https://cran.r-project.org/package=ReviewR) and is distributed with the open source 3-Clause BSD License on GitHub (https://github.com/thewileylab/ReviewR/). Users can also trial ReviewR without download using ShinyApps.io (https://thewileylab.shinyapps.io/ReviewR/). All versions include a SQLite demonstration database that includes 10 records from the CMS SynPUF data set and "clinical notes" from a medical transcription training corpus.[46–48]

## Using ReviewR

Users access the Setup tab to connect ReviewR to their patient database (Figure 1). Users select the RDBMS from the drop down menu which opens the connection module. Completing this connection sets the user into *view mode* where they can explore patient records. Optionally, users can enter *review mode* that supports data abstraction by connecting to and configuring REDCap. Users must have created a REDCap project and instrument through their institution and have access to an API key for the project. Once connected, users are prompted to configure the instrument which allows ReviewR to automatically populate both the patient and reviewer names in the instrument, reducing errors and repetitive data entry tasks. After connecting, users select records to review from the "Patient Search" tab. This tab displays a list of patients, demographics, and abstraction status (if configured). Users can click on the patient identifier to navigate to the individual's record.

Individual patient records are displayed in the "Chart Review" tab (Figure 2). Patient demographic information is always available for the selected patient at the top of the screen, along with navigation options to progress through the patients. Below the demographic data is a tabbed interface containing the information from each table in the connected database. Users can easily search structured and unstructured data using regular expressions both within and across columns, tabs/tables, and even records, filtering to only the relevant entries with matches highlighted for easy identification.

When abstraction is configured, the REDCap abstraction instrument will be displayed to the right of the patient data. Data are uploaded to REDCap using the "Save to REDCap" button, with ReviewR warning users if it detects that previously entered data have been changed.

## Extending ReviewR

We have included a number of developer functions to help users extend ReviewR to support additional data models and RDBMSs. Adding a new data model requires limited R programming experience and is outlined in Figure 3. Users provide a data model schema as a CSV file and then select the table and fields containing patient information. More technical users may optionally customize table displays by editing the ReviewR generated R code. Adding a new RDBMS requires more programming expertise as users must develop a Shiny module that captures the database credentials and returns a valid {DBI} connection. However, ReviewR offers functions and templates to support this process. Vignettes describing this process and full package documentation are available at https://reviewr.thewileylab.org.

## Demonstration projects

The first demonstration project was performed at Northwestern University which restricts EHR access for clinical use only. Researchers needed to perform chart review to confirm the case/control status of a random subset of patients identified by a heart failure phenotyping algorithm. LVR successfully built a {DBI} connection object to connect ReviewR to the Northwestern Medicine Enterprise Data Warehouse OMOP instance which is stored in Microsoft SQL Server. ReviewR was then used to perform a technical evaluation of heart failure phenotyping algorithms; the results of which were subsequently published.[49]

The second demonstration project was performed at the University of Colorado to test the face validity of chart review conducted



**Figure 1.** ReviewR setup tab. This page allows users to enter the view mode by connecting to their patient database first selecting the appropriate RDBMS and then providing connection credentials (for Postgres) or using the "Sign in to Google" interface (for Bigquery). Optionally, users can enter review mode by connecting to a previously created REDCap project using an API key and then configuring the connection to identify the review field that will hold the patient and (optionally) reviewer identifiers.

**Figure 2.** ReviewR chart review tab. This page allows users to review the medical chart for a single patient. The patient information is shown at the top, including the chart review status as recorded in REDCap. Chart information shows all of the available tables for the configured data model. Users can perform a search (including the use of regular expressions) across all of the tables, and results will be highlighted (A). For each of the columns in a data table, users can do additional filtering and sorting (B). Users are guided through charts using quick navigation controls, and the configured REDCap instrument is displayed alongside the chart data.

in ReviewR compared to the source EHR and to build a gold-standard cohort for aneurysm phenotype algorithm development. Although 50 records were identified for review, one record was accidentally skipped during ReviewR extraction and was excluded from our analysis. There was an overall concordance rate of 94% (Table 2), with ReviewR identifying 2 patients with a history of an aneurysm and 1 patient with a history of SAH that were not identified during chart review in Epic. Of these, 1 patient only had documentation of the aneurysm in outside media records that were not viewed during the original Epic-based search. The other 2 discordant patients' clinical notes documented an aneurysm/SAH, but brain imaging showed no clips/coils present—contradicting the record.

In addition to these known uses, ReviewR has been downloaded from CRAN 2315 times between April 2, 2021 and October 16, 2021. Unfortunately, further details about the identities or of the downloaders or their planned use cases are not available.

## DISCUSSION

To our knowledge, ReviewR represents the first open source tool explicitly designed to support manual record review that provides a unified environment for data review and chart abstraction data entry. ReviewR was designed to meet multiple research challenges that we have observed within our research experience. First and foremost, the unified interface provides researchers with a streamlined workflow while automated entry of patient and reviewer identifiers target common sources of data entry errors and inefficiencies. Second, ReviewR fills gaps present at many institutions where research access to source EHRs is restricted or researchers only have access to deidentified data warehouses. In many cases these institutions also lack record-level display tools making manual record review challenging and laborious. R Shiny provides a lightweight design sup-

porting standalone deployment by individual researchers who may lack technical infrastructure and support staff to operate more robust tools.

A key strength of ReviewR comes from its extensible design and use of Shiny modules. By partitioning the application into a modular structure, adding support for a new data model or RDBMS requires little to no modification of ReviewR core code. Further the inclusion of developer functions to streamline, and extensive documentation to support, this process will hopefully attract a community of users who can help grow the functionality of ReviewR. Additionally, these modules can be easily repurposed by other researchers in their own applications. For example, the database modules all provide generic connection objects that are broadly useful.

Finally, our demonstration projects prove that ReviewR can be deployed in practice to support chart review across institutions and use cases. The 2 projects used different RDMSs and data models validating the extensibility and flexibility of the tool. Moreover, our paired testing of ReviewR compared to Epic-based chart review helps support the face validity of performing chart review using clinical data warehouses. Although these warehouses typically contain only a portion of the patient record, in our testing ReviewR reveal high concordance with the source EHR text and structured data. The comprehensive search capabilities of ReviewR revealed an additional case that was missed during Epic review due to supporting information only being located in ancillary documents. Two false-positive ReviewR findings were only detectable using clinical images, representing a fundamental limitation of chart review performed on clinical data warehouses that do not contain imaging data.

However, ReviewR does have limitations. First, the user interface for ReviewR has not been formally evaluated or studied. Although anecdotally our experience has been positive, a more formal assessment is needed before any claims could be made about gains in efficiency. Second, ReviewR has a limited number of databases and
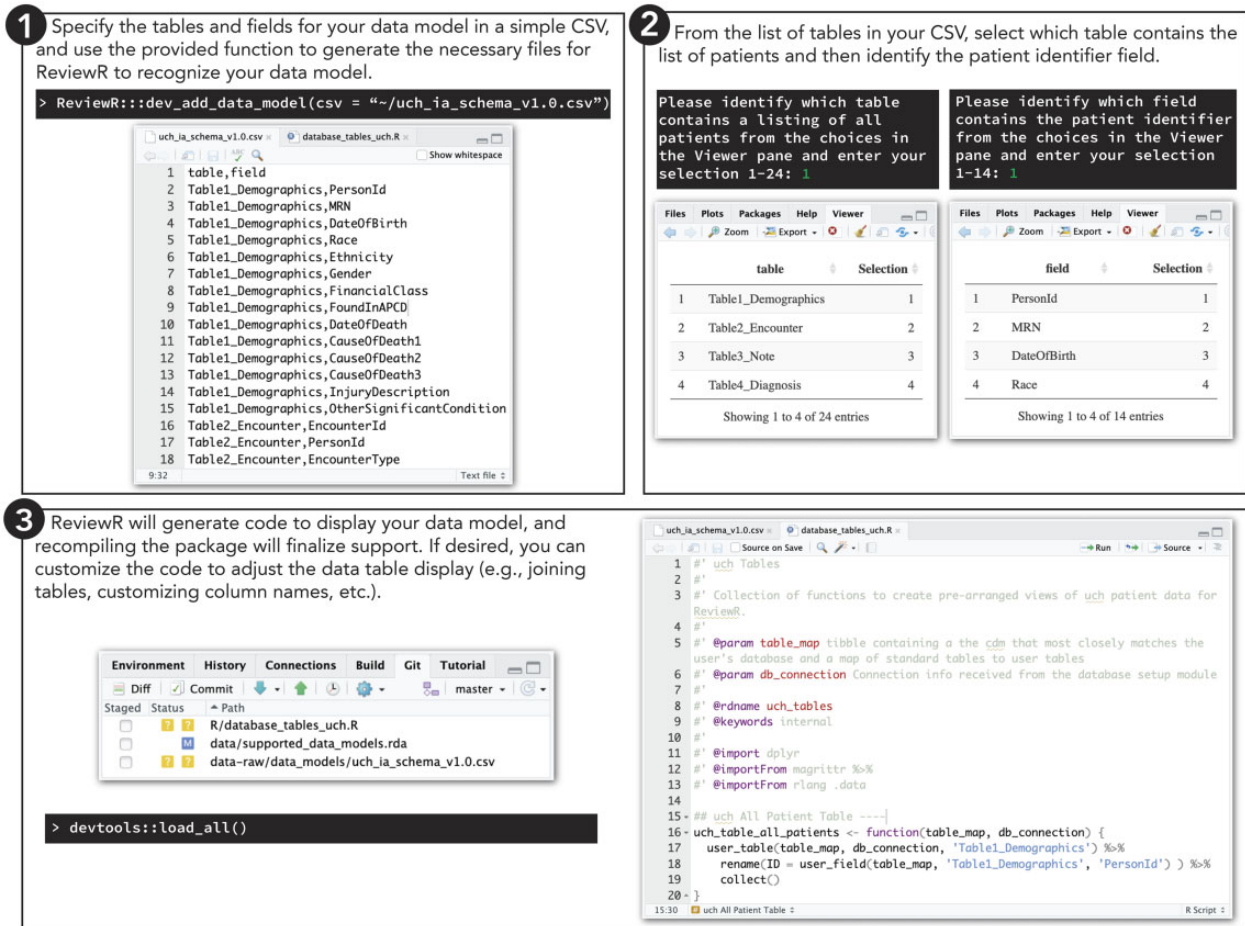
**Figure 3.** ReviewR process for adding support for a custom data model. Users can easily add support for a new data model using build in developer tools. Step 1: The user must create a schema file that contains all tables and associated field names as a csv file. Within the R console they pass this file to the "dev_add_data_model()" function. Step 2: ReviewR will walk users through selecting the table containing patient demographics (used to define the "Patient Search" tab) and the column name containing the patient identifier (used to auto populate the REDCap instrument). Step 3: ReviewR adds the schema to the list of supported data models and generates template R code for displaying the database as is. Reloading the package finalizes this support. More technical users may customize the R code to change table displays, join tables (eg, the OMOP data model requires joining to the concept table to have an informative information display), etc.

**Table 2.** Second demonstration project review results

| Extraction order | Concordant | Discordant | Total |
|---|---|---|---|
| ReviewR -> Epic | 23 | 2[a] | 25 |
| Epic -> ReviewR | 23 | 1[b] | 24 |
| Total | 46 | 3 | 49 |

   [a]ReviewR identified 1 additional aneurysm and aneurysmal SAH compared to Epic review.

   [b]ReviewR identified 1 additional aneurysm compared to Epic review.

data models that it currently supports, and only connects to data warehouses and not EHR systems (though with the development of R packages to support FHIR-based exchange this may become possible). Third, ReviewR requires clinical database access, however, while many institutions restrict access to the core data warehouse it is increasingly common for clinical data to be delivered to researchers using database systems.[50] Finally, although R is widely adopted, the selection of R Shiny as the application platform for ReviewR may preclude its deployment as an enterprise-wide solution where R Shiny is not currently supported. By allowing ReviewR to run in multiple modes (standalone, hosted, Docker container), we hope that potential users will find a secure deployment option that works for their organization.

## CONCLUSION

In this paper, we describe the development and implementation of ReviewR—a flexible and extensible tool that can be used to perform chart abstractions from EHR data stored in clinical data warehouses. Its unified user interface streamlines the researcher workflow and can reduce potential errors during the review process. As a free and open source solution, we hope its continued adoption and refinement can improve the speed and quality of manual record review.

## FUNDING

## AUTHOR CONTRIBUTIONS

LKW conceived the idea of ReviewR, acquired funding, and supervised the overall direction and scope of the project. LKW and LVR developed an initial prototype of the software; DAM developed the release version of ReviewR and drafted all software documentation. LVR and CDR performed beta testing, and MGK and LMS provided substantive input on design and functionality of ReviewR. LKW and DAM developed the initial draft of this manuscript, and all authors reviewed and approved the final manuscript.

## ACKNOWLEDGMENTS

## CONFLICT OF INTEREST STATEMENT

LKW and MGK have received funding from Google Cloud Healthcare to support education through the Coursera Clinical Data Science Specialization and through that partnership received helpful technical advice on ReviewR's support of Google BigQuery.

## DATA AVAILABILITY

Data from the second demonstration project are openly available on Dryad at https://doi.org/10.5061/dryad.bk3j9kdfx

## REFERENCES

1. Nordo AH, Levaux HP, Becnel LB, *et al*. Use of EHRs data for clinical research: historical progress and current applications. *Learn Health Syst* 2019; 3 (1): e10076.
2. Bowton E, Field JR, Wang S, *et al*. Biobanks and electronic medical records: enabling cost-effective research. *Sci Transl Med* 2014; 6 (234): 234cm3.
3. Casey JA, Schwartz BS, Stewart WF, Adler NE. Using electronic health records for population health research: a review of methods and applications. *Annu Rev Public Health* 2016; 37: 61–81.
4. Sarkar S, Seshadri D. Conducting record review studies in clinical practice. *J Clin Diagn Res* 2014; 8 (9): JG01–4.
5. Newton KM, Peissig PL, Kho AN, *et al*. Validation of electronic medical record-based phenotyping algorithms: results and lessons learned from the eMERGE network. *J Am Med Inform Assoc* 2013; 20 (e1): e147–54.
6. Alzu'bi AA, Watzlaf VJM, Sheridan P. Electronic health record (EHR) abstraction. *Perspect Health Inf Manag* 2021; 18 (Spring): 1g.

7. Brundin-Mather R, Soo A, Zuege DJ, *et al*. Secondary EMR data for quality improvement and research: a comparison of manual and electronic data collection from an integrated critical care electronic medical record system. *J Crit Care* 2018; 47: 295–301.
8. Harris PA, Taylor R, Thielke R, *et al*. Research electronic data capture (REDCap)–a metadata-driven methodology and workflow process for providing translational research informatics support. *J Biomed Inform* 2009; 42 (2): 377–81.
9. Murphy SN, Weber G, Mendis M, *et al*. Serving the enterprise and beyond with informatics for integrating biology and the bedside (i2b2). *J Am Med Inform Assoc* 2010; 17 (2): 124–30.
10. Atlas [Software]. (GitHub). https://github.com/OHDSI/Atlas.
11. Dobbins NJ, Spital CH, Black RA, *et al*. Leaf: an open-source, model-agnostic, data-driven web application for cohort discovery and translational biomedical research. *J Am Med Inform Assoc* 2020; 27 (1): 109–18.
12. Hanauer DA, Mei Q, Law J, Khanna R, Zheng K. Supporting information retrieval from electronic health records: a report of University of Michigan's nine-year experience in developing and using the Electronic Medical Record Search Engine (EMERSE). *J Biomed Inform* 2015; 55: 290–300.
13. Soysal E, Wang J, Jiang M, *et al*. CLAMP—a toolkit for efficiently building customized clinical natural language processing pipelines. *J Am Med Inform Assoc* 2018; 25 (3): 331–6.
14. Ogren PV. Knowtator: a protégé plug-in for annotated corpus construction. In: Proceedings of the 2006 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology Companion Volume: Demonstrations. Association for Computational Linguistics, 2006. doi:10.3115/1225785.1225791.
15. Stenetorp P, Pyysalo S, Topić G, Ohta T, Ananiadou S, Tsujii J. brat: a web-based tool for NLP-assisted text annotation. In: Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics. Avignon, France: Association for Computational Linguistics; 2012: 102–7.
16. Chang W, Cheng J, Allaire JJ, Xie Y, McPherson J. *shiny: Web Application Framework for R [Software]*. 2017. https://CRAN.R-project.org/package=REDCapR.
17. Guyader V, Fay C, Rochette S, Girard C. *golem: A Framework for Robust Shiny Applications [Software]*. 2020. https://github.com/ThinkR-open/golem.
18. Fay C, Guyader V, Rochette S, Cervan G. *Engineering Production-Grade Shiny Apps*. New York: CRC Press; 2021.
19. R Core Team. *R: A Language and Environment for Statistical Computing [Software]*. 2013. https://www.R-project.org/.
20. Wickham H, Hester J, Chang W. *devtools: Tools to Make Developing R Packages Easier [Software]*. 2020. https://CRAN.R-project.org/package=devtools.
21. Wickham H, Averick M, Bryan J, *et al*. Welcome to the Tidyverse. *J Open Source Softw* 2019; 4 (43): 1686.
22. Allaire JJ. *config: Manage Environment Specific Configuration Values [Software]*. 2020. https://CRAN.R-project.org/package=config.
23. Müller K. *here: A Simpler Way to Find Your Files [Software]*. 2020. https://CRAN.R-project.org/package=here
24. Cheng J, Sievert C, Chang W, Xie Y, Allen J. *htmltools: Tools for HTML [Software]*. 2021. https://CRAN.R-project.org/package=htmltools.
25. Csárdi G, Chang W. *processx: Execute and Control System Processes [Software]*. 2020. https://CRAN.R-project.org/package=processx.
26. Hester J. *glue: Interpreted String Literals [Software]*. 2020. https://CRAN.R-project.org/package=glue
27. Grosser M. *snakecase: Convert Strings into Any Case [Software]*. 2019. https://CRAN.R-project.org/package=snakecase.
28. Lilovski N. *dashboardthemes: Customise the Appearance of 'shinydashboard' Applications Using Themes [Software]*. 2020. https://CRAN.R-project.org/package=dashboardthemes.
29. Chang W, Borges Ribeiro B. *shinydashboard: Create Dashboards with 'Shiny' [Software]*. 2018. https://CRAN.R-project.org/package=shinydashboard.

30. Granjon D. *shinydashboardPlus: Add More 'AdminLTE2' Components to 'shinydashboard' [Software]*. 2021. https://CRAN.R-project.org/package=shinydashboardPlus.

31. Perrier V, Meyer F, Granjon D. *shinyWidgets: Custom Inputs Widgets for Shiny [Software]*. 2021. https://CRAN.R-project.org/package=shinyWidgets.

32. Attali D. *shinyjs: Easily Improve the User Experience of Your Shiny Apps in Seconds [Software]*. 2020. https://CRAN.R-project.org/package=shinyjs.

33. Sali A, Attali D. *shinycssloaders: Add Loading Animations to a 'shiny' Output While It's Recalculating [Software]*. 2020. https://CRAN.R-project.org/package=shinycssloaders.

34. Xie Y. *DT: A Wrapper of the JavaScript Library 'DataTables' [Software]*. 2018. https://CRAN.R-project.org/package=DT.

35. DataTables [Software]. https://datatables.net/.

36. Wickham H, Girlich M, Ruiz E. *dbplyr: A 'dplyr' Back End for Databases [Software]*. 2021. https://CRAN.R-project.org/package=dbplyr.

37. R Special Interest Group on Databases (R-SIG-DB), Wickham, H. & Müller, K. DBI: R Database Interface [Software]. 2021. https://CRAN.R-project.org/package=DBI.

38. Müller K, Wickham H, James DA, Falcon S. *RSQLite: 'SQLite' Interface for R [Software]*. 2021. https://CRAN.R-project.org/package=RSQLite.

39. Wickham H, Ooms J, Müller K. *RPostgres: 'Rcpp' Interface to 'PostgreSQL' [Software]*. 2021. https://CRAN.R-project.org/package=RPostgres.

40. Wickham H, Bryan J. *bigrquery: An Interface to Google's 'BigQuery' 'API' [Software]*. 2020. https://CRAN.R-project.org/package=bigrquery.

41. Bryan J, Citro C, Wickham H. *gargle: Utilities for Working with Google APIs [Software]*. 2021. https://CRAN.R-project.org/package=gargle.

42. Hripcsak G, Duke JD, Shah NH, *et al*. Observational Health Data Sciences and Informatics (OHDSI): opportunities for observational researchers. *Stud Health Technol Inform* 2015; 216: 574–8.

43. Johnson A, Pollard T, Mark R. MIMIC-III Clinical Database Demo. *PhysioNet* 2019. doi:10.13026/C2HM2Q.

44. Beasley W. *REDCapR: Interaction between R and REDCap [Software]*. 2020. https://CRAN.R-project.org/package=REDCapR.

45. Nutter B, Lane S. *redcapAPI: Accessing Data from REDCap Projects Using the API [Software]*. 2020. doi:10.5281/zenodo.11826.

46. CMS 2008-2010 Data Entrepreneurs' Synthetic Public Use File (DE-SynPUF). https://www.cms.gov/Research-Statistics-Data-and-Systems/Downloadable-Public-Use-Files/SynPUFs/DE_Syn_PUF. Accessed October 5, 2021.

47. ETL-CMS: Workproducts to ETL CMS Datasets into OMOP Common Data Model [Software]. (GitHub). https://github.com/OHDSI/ETL-CMS.

48. Transcribed Medical Transcription Sample Reports and Examples—MTSamples. https://www.mtsamples.com/. Accessed October 5, 2021.

49. Brandt PS, Kiefer RC, Pacheco JA, *et al*. Toward cross-platform electronic health record-driven phenotyping using clinical quality language. *Learn Health Syst* 2020; 4 (4): e10233.

50. Kahn MG, Mui JY, Ames MJ, *et al*. Migrating a research data warehouse to a public cloud: challenges and opportunities. *J Am Med Inform Assoc* 2021. http://dx.doi.org/10.1093/jamia/ocab278.