Paper DV09

Empowering Clinical Data Review with Interactive Patient-Level Visualisations in {teal}

Dony Unardi, Roche/Genentech, South San Francisco, USA
Nina Qi, Roche/Genentech, South San Francisco, USA

ABSTRACT

The release of {teal} v1.0.0 on CRAN marks a major milestone in the evolution of this open-source framework for interactive clinical data exploration. This presentation focuses on significant feature enhancements that extend {teal}'s flexibility and scalability across the pharmaceutical industry. Key among these is the introduction of new {teal} modules that enable interactive patient-level drill-downs. By integrating Plotly, we have developed specific modules to visualise safety and efficacy data, allowing reviewers to filter, explore, and analyze subject-level information in real time. This capability enhances transparency and efficiency in medical data review. Furthermore, the v1.0.0 release includes a redesigned user experience with streamlined navigation and modernized layouts, as well as a refined {teal.reporter} for standardized, reproducible reporting, reinforcing {teal}'s adoption and usability.

INTRODUCTION

Clinical data review plays a central role in evaluating safety and efficacy throughout drug development. While traditional review relies on static tables, listings, and figures, these outputs can limit a reviewer's ability to explore data dynamically, investigate emerging signals, and efficiently trace findings back to individual patients. As clinical trials increase in complexity and data volume, reviewers need tools that support flexible exploration, rapid iteration, and seamless movement between population-level summaries and patient-level detail.

Interactive visualisation provides a powerful complement to static clinical reporting by enabling real-time data exploration and patient-level drill downs. The {teal} framework addresses this need through an open-source, modular platform for interactive clinical data review built on R and Shiny. With the release of {teal} version 1.0.0 on CRAN, the framework has reached a key milestone in maturity and usability. This paper presents a key component of the recent enhancements to {teal} that introduce interactive patient-level visualisations, illustrating how these capabilities can support more efficient, transparent, and insightful clinical data review.

OVERVIEW OF THE {teal} FRAMEWORK

The {teal} framework is an open-source platform for building interactive clinical data exploration applications using R. It follows a modular design, where individual analyses, visualisations, or tables are packaged as teal modules. These modules can be combined within a single application, allowing teams to assemble interactive review tools in a consistent and flexible way. The framework provides a ready-to-use application structure and includes essential features such as data filtering, reproducible R code generation, and reporting, so developers can focus more on analysis content rather than application setup.

A typical {teal} application presents results in a display area, supported by user controls for filtering data and adjusting analysis settings. Users can move easily between different analysis modules, explore results interactively, and examine how outputs change as inputs or parameters are modified. To support reuse and extension, {teal} is part of a broader ecosystem of packages, which provides curated collections of modules for clinical trial, biomarker, and general exploratory analyses.

The release of {teal} version 1.0.0 on CRAN marked an important milestone, providing a stable and well-defined foundation for wider use. This release introduced a refreshed user interface with clearer navigation and updated layouts, improving the overall user experience. Subsequent updates in version 1.1 further enhanced the framework by integrating an improved {teal.reporter}, enabling users to capture interactive analyses and outputs into standardized and reproducible reports. These developments provide a strong platform for extending {teal} with more advanced capabilities, including interactive patient-level visualisations.

INTERACTIVE PATIENT-LEVEL VISUALISATIONS

Motivation

In many clinical development settings, patient-level data review relies on established and centrally managed platforms that provide standardized and consistent analysis outputs. These platforms play an important role in supporting routine review activities. However, as clinical trial designs become increasingly complex and adaptive, additional study-specific clinical questions often emerge throughout the course of a study. Addressing these evolving questions may require analyses or visualisations that extend beyond predefined configurations. A key limitation of existing solutions is not interactivity itself, but the difficulty of modifying and customizing analyses to address evolving, study-specific clinical questions, particularly when such modifications require additional time and specialized resources.

This lack of flexibility has practical implications for clinical scientists, who are the primary users of patient-level data during safety and efficacy evaluation. When exploratory review is constrained, clinical scientists may need to perform additional manual review steps, ad-hoc exploration, or offline analyses to address emerging clinical questions. This can reduce the time available for interpretation, insight generation, and decision-making, and can introduce inefficiencies into the overall review workflow.

To address these challenges, interactive patient-level visualisations were developed within the {teal} framework. By providing an open, modular, and accessible platform for exploratory clinical data review, {teal} is intended to complement existing standardized processes while enabling clinical scientists to explore patient-level data more flexibly and efficiently, with improved transparency and traceability during review.

Design Approach: Enabling Patient-Level Drill Downs

To support flexible and interactive patient-level exploration, interactive visualisation was adopted as a core design principle. This approach focuses on enabling users to move seamlessly between summary-level views and subject-level detail while maintaining context throughout the review process. To achieve this, we leveraged established open-source R packages that are well suited for interactive clinical data review, including plotly, reactable, and bslib.

The plotly R package was used to create interactive, web-based graphics built on the open-source JavaScript library plotly.js. When combined with Shiny, Plotly visualisations support real-time interaction between visual elements and underlying data. A key capability in this setting is server-side linking, which allows user interactions such as selecting data points or filtering data to trigger updates in other views. This enables intuitive drill-down workflows, where users can begin with aggregated summaries and progressively explore patient-level detail without leaving the application or requesting additional outputs.

To complement graphical exploration, the reactable package was used to support interactive tabular displays that respond dynamically to user selections. This allows patient-level tables to remain synchronized with visual selections, providing immediate access to detailed subject information as users explore signals of interest. In addition, the bslib package was used to support consistent and modern user interface styling, improving usability and visual clarity across interactive components.

Within the {teal} framework, these capabilities are embedded directly into modular analyses. Clinical scientists can explore safety and efficacy patterns dynamically, trace observations back to individual patients, and iteratively refine their focus in real time. By emphasizing linked interactivity across graphical and tabular views, rather than static outputs, this design approach supports a more efficient, transparent, and user-driven review process that aligns with evolving clinical data review needs.

Use Cases and Early Adoption

Interactive patient-level modules developed within the {teal} framework have been piloted across multiple clinical studies spanning different therapeutic areas. These modules focus on common safety and efficacy review needs, enabling users to move seamlessly from summary-level views to subject-level detail through interactive drill downs. While real study data were used during development and evaluation, visual examples are not included in this paper due to data confidentiality considerations.

Feedback sessions conducted with clinical scientists demonstrated strong interest and positive reception. Users highlighted the ability to explore data dynamically, trace observations back to individual patients, and adjust their focus in real time as key advantages over traditional static outputs. Importantly, these capabilities allowed clinical scientists to spend less time on manual data checks and ad-hoc reporting, and more time interpreting results and supporting clinical decision-making. These early pilots suggest that interactive patient-level visualisations can meaningfully enhance the efficiency and effectiveness of medical data review.

Operational Consideration: Validation

Validation is an important consideration for interactive tools used in regulated environments such as medical data review. For this initiative, the {teal} framework and the associated patient-level modules have undergone internal software validation to support their intended use within clinical development workflows.

In addition to framework-level validation, each study-specific Shiny application built using {teal} follows a separate validation process to ensure that the final application is appropriate for the specific study context and data. This separation allows the underlying framework and reusable components to remain stable, while enabling flexibility to address study-specific review needs.

This layered validation approach supports the compliant use of interactive patient-level visualisations as part of routine medical data review. It also enables interactive tools to be applied in a controlled manner while accommodating the evolving nature of clinical data and exploratory review activities.

CONCLUSION

Interactive patient-level data review is becoming increasingly important as clinical trials grow in complexity and data volume. Traditional static outputs and rigid review tools can limit flexibility, slow down insight generation, and place unnecessary operational burden on both programming teams and clinical scientists. Addressing these challenges requires solutions that are not only analytically powerful, but also accessible, scalable, and fit for use in regulated environments.

This paper presented recent enhancements to the {teal} framework that enable interactive patient-level visualisations for clinical data review. By combining {teal}'s modular architecture with linked interactive graphics, users can explore safety and efficacy data dynamically, trace findings back to individual patients, and adapt their review focus in real time. Early piloting across multiple studies and therapeutic areas, along with positive feedback from clinical scientists, demonstrates the potential of this approach to improve efficiency, transparency, and decision support in medical data review.

As adoption of open-source technologies continues to expand in clinical development, frameworks such as {teal} provide a sustainable foundation for modern data review workflows. Future work will focus on further expanding patient-level modules, refining usability based on user feedback, and supporting broader adoption of interactive, reproducible clinical data exploration across studies and organizations.

ACKNOWLEDGMENTS

The authors would like to thank the clinical scientists who participated in feedback sessions and shared their insights during the piloting of the interactive patient-level visualisations described in this paper. We also acknowledge the contributions of colleagues who supported the development, maintenance, and application of the {teal} framework and its associated modules. In addition, we recognize the broader R and open-source communities, including the developers and maintainers of R/Shiny, plotly, reactable, and bslib, whose work enables interactive and flexible clinical data exploration.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Author Name: Dony Unardi
Company: Roche/Genentech
Email: unardi.dony@gene.com

Co-Author Name: Nina Qi
Company: Roche/Genentech
Email: qi.ting@gene.com