Paper CM06

An Interactive Visualization Application for Patient Reported Outcomes with Spotfire

Matt Buckley, Hoffmann-La Roche Limited, Mississauga, Canada

ABSTRACT

Interactive visualization tools provide researchers with a quick way to answer pertinent research questions and guide further analysis. These tools are especially helpful in the domain of Patient Reported Outcomes. This work demonstrates a TIBCO Spotfire™ application built for this domain. It highlights demographic data, clinically meaningful changes, questionnaire completion rates, and assessment measures by visit. It is also highly customizable, allowing users to filter for the population of interest and segment analyses by stratification variables. However, for large datasets and detailed analyses, Spotfire can become unresponsive and slow. The application detailed in this work was optimized for efficient calculations and low storage costs by pre-processing data in SAS. This makes the application responsive and quickly configurable. In summary, Spotfire provides a flexible and powerful framework for interactive visuals, and with some efficiency considerations, is able to handle complex analyses and provide key insights.

1. BACKGROUND

Spotfire is a powerful tool for interactive visualizations in the form of a dashboard. This allows researchers to interpret large quantities of data and form hypotheses rapidly. The interactive nature allows investigators to filter views for populations of interest and focus their attention on the most important information. Then, with the help of more thorough analyses, these hypotheses can be investigated further.

2. INTRODUCTION

This work discusses a Spotfire application built to visualize a subset of clinical data known as Patient Reported Outcomes (PROs). In this domain, patients complete a series of questionnaires at various visits throughout a clinical trial. The aim is to assess how patients feel and act[1]. Spotfire, has long been used as a tool within Roche to visualize clinical trial data for the purposes of supplementary analysis. Roche has already used Spotfire for PRO data in other clinical trials. However, previous PRO templates were not built on the CDISC Analysis Data Model (ADaM)[2], thus a new application was needed to work with ADaM compliant datasets.

Within Roche, the Patient Centered Outcomes Research (PCOR) department works with and interprets PRO data in the context of clinical trials. Often, there is a very short time frame between when a clinical trial is unblinded and when the writing of the primary clinical study report finishes. Tables, listings, and graphs need to be regenerated and checked for accuracy within this timeframe. A Spotfire dashboard allows PCOR scientists to see the conduct preliminary analysis before seeing the official results. In addition, these scientists can perform additional analyses not covered in the tables, listings, and graphs.

However, for the tool to be effective it needs to be quick to set up and easily adaptable. Efficiency and reusability were the main goals when designing this tool. Spotfire's main strength is in data visualization and so most of the pre-processing and data set up was instead performed in SAS. SAS is a robust statistical language capable of dynamic and efficient processing. All operations are performed in a single script, adhering to the DRY (Don't Repeat Yourself) coding principle[3]. This makes it simple for others to modify the application to fit their needs while also offering performance gains. The SAS script performs all complex calculations one time, reducing redundancy. If instead the operations had been performed in Spotfire, then these calculations would be recalculated each time a user opened the application or switched between visuals.

3. APPLICATION LAYOUT

The application displays patient demographic data, a summary of clinically meaning changes, questionnaire completion rates, and assessment measures as both absolute values and change from baseline. In addition, the application displays information relating to disease burden and time to deterioration data. The analyses are grouped by visit and the user is able to specify additional grouping variables such as treatment arm, sex, and race. Users can easily change the population studied through a selection box, enabling views of groups such as safety population, intent-to-treat population, or PRO-evaluable population. Furthermore, users can alter thresholds and grouping variables in order to generate analyses that may not have been planned at study design.

The selection tools are placed on a primary setup page, as depicted in figure 1. This page displays summary level data and a brief guide detailing how to use the application. When a user selects a filter on this page, the filter applies throughout the entire application. However, this filter then appears on each page to which it applies. For instance, if the intent-to-treat population is selected, then this selection is displayed on every page. This allows the user to create a consistent story throughout the application without losing track of the current filters applied.

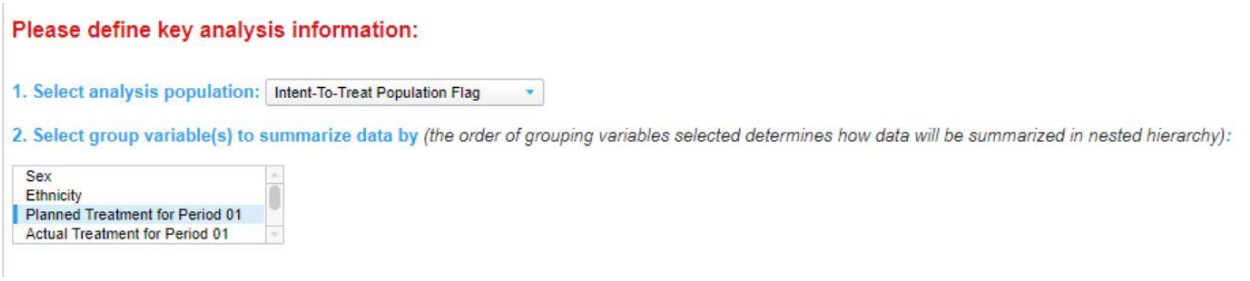

Figure 1: Primary Setup Selections

Spotfire also offers the ability to assign certain variables as parameters. This functionality is used in the application to ensure that separate charts and graphs are all grouped by the same variables. This procedure can also be combined with the filtering and marking capabilities of Spotfire. Marking allows the user to highlight a certain subgroup of data and then this subgroup is highlighted in all related visuals throughout the application. By combining these capabilities, scientists can analyze certain subgroups of patients throughout the clinical trial. In addition, it allows scientists to form and investigate new hypotheses in a matter of minutes. With traditional tables, listings, and graphs, data supporting these new hypotheses would have to be coerced into useful outputs. If it then turned out that the hypotheses were incorrect, this work would have been wasted.

Visualization tools such as Spotfire are not meant to replace the traditional outputs sent to regulatory bodies such as the FDA. Instead, they allow other groups earlier access to the data and empower these individuals to explore new ideas and answer important research questions. These researchers and scientists can then form new stories about a particular group of patients or the trial as a whole. Then, if some of these investigations lead to insights that are important to the trial, these can be included in the official outputs to provide internal stakeholders and regulatory authorities with a more complete picture.

4. DEVELOPMENT

The application was built in an agile manner. The development process began with a minimal version of the application. Then feedback was continuously gathered throughout the development process from various stakeholders. This ensures that the application is user friendly and meets the needs of the users. It also encourages users to ask questions throughout the development process. These questions can be for the purpose of understanding or they may lead to new features not thought of by the original developer. In addition, the application was built as a template. The template can be easily tailored to meet the needs of other clinical studies while automating a large portion of redundant work.

Figure 2 depicts the continuous and ongoing nature of the development process. The developer constantly needs to either add new features or modify existing functionality based on user feedback. In this way, the design of the application becomes a team effort and ensures user requirements are met.

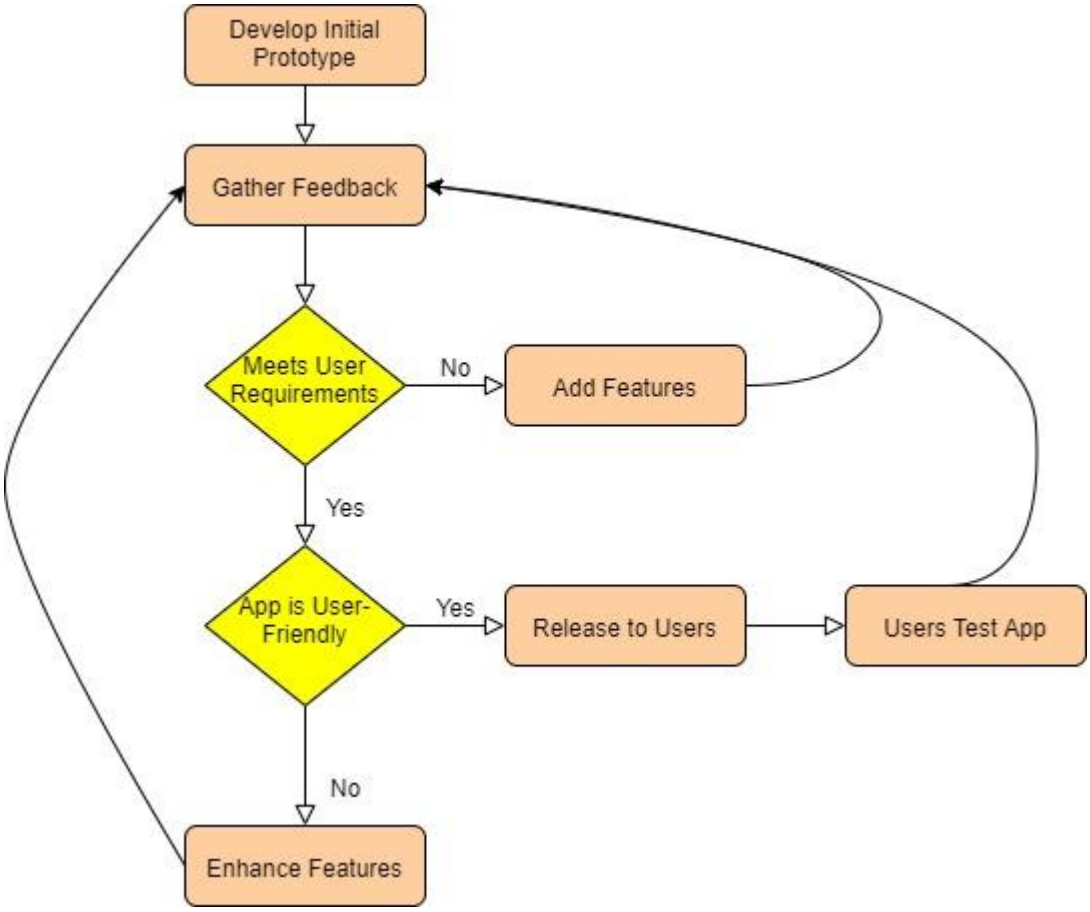

Figure 2: The Development Workflow

While Spotfire provides a strong analytic and data visualization framework, for large analyses with numerous calculations it can become slow and unresponsive. Some calculations are completed efficiently by Spotfire and the results are stored in memory. For others, the results are computed each time the data is loaded. In large applications where data is accessed outside of Spotfire, this leads to significant delays in the responsiveness and diminishes the overall quality of the user experience. As a result, most of the data manipulations and transformations needed to generate these visuals are done in SAS.

The pre-processing is completed in a single SAS script. By applying a series of transformations, the script transforms the input ADaM datasets into a new set of datasets that are ready to be used by Spotfire. The script is organized in small manageable sections that are both easy to follow and to change. In addition, two datasets of metadata are created. The first specifies characteristics of the input ADaM data while the second includes similar characteristics for the data created by the script, and which appear in the Spotfire dashboard. This adds traceability, allowing users to see where the data came from and which transformations were applied to generate the report. Figure 3 depicts the flow of data which leads to the application generation.

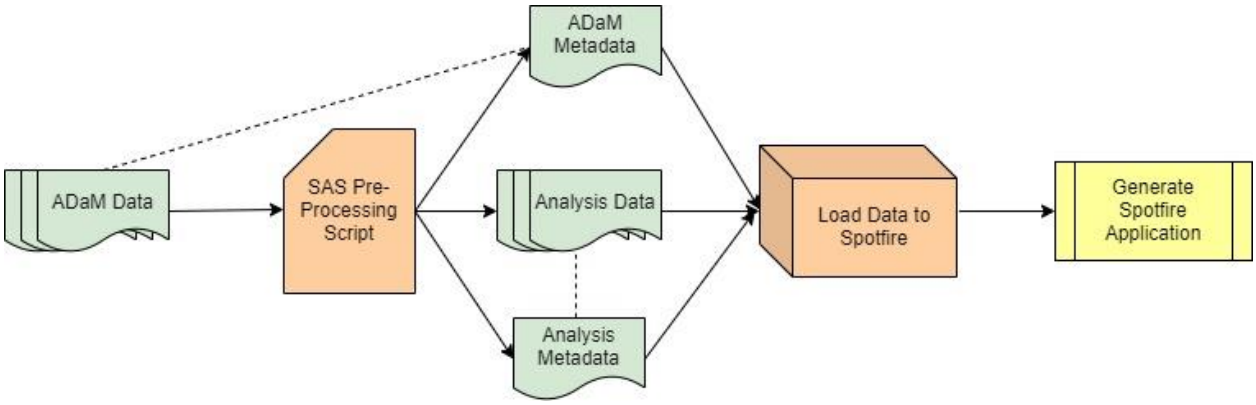

Figure 3: Development Data Flow

5. REUSE

The original Spotfire application explained in this work was developed using the EORTC QLQ-C30 Questionnaire[4]. It is a questionnaire designed by the European Organization for Research and Treatment of Cancer (EORTC), used to assess the quality of life for cancer patients. It consists of 30 questions, answered by the patients. These answers are then combined into a set of measures/scales meant to assess various aspects of how patients feel and what they experience. Some examples of scales that are measured include Global Health Status, Physical Functioning, and Role Functioning.

Although the application was initially developed using the EORTC QLQ-C30 questionnaire, these details are only specified in the ADaM data and the SAS pre-processing script. This can easily be altered for other questionnaires by making small changes to the SAS script. As explained above, the Spotfire template uses parameters to reference variables used in the charts and ensure consistency across the visuals. When setting up the application, a user can easily change these parameters to use a different set of variables. There is also one section of the pre-processing script that references the set of key variables used in the application. These can then be changed to reflect alternate needs of different studies.

The set of key variables defined in the SAS script play a crucial role throughout the application and serve as the main variables for most visuals. These are defined as new variables and merged with the existing ADaM datasets. By defining these variables separately, two goals are accomplished. First, all key variable definitions can be found in a single place. This makes it easy to find and modify. Second, even within the ADaM data model, there is flexibility to accommodate different study needs. As a result, not every study has the same set of variables. These new variables allow the template to be used across studies with only minor modifications required to the SAS script.

Multiple oncology studies at varying stages have successfully used this template. The SAS script can be modified to fit the needs of each individual study. Once set up, the application can be updated with new data quickly. For one study, the application was set up and maintained during the blinded phase. Once the PCOR scientists had been unblinded, the application was modified with the unblinded data and sent to the scientists later that same day. This allowed the scientists to view and analyze the data long before they normally would have received the study outputs. As a result, they were able to successfully perform analyses and investigate key questions under tight deadlines.

6. CONCLUSION

Spotfire is a powerful data visualization tool. Roche utilizes Spotfire successfully as a supportive tool for clinical studies, guiding the analysis of researchers and allowing these users to investigate new hypotheses as they arise. This work demonstrates a situation in which Spotfire can successfully be used for supportive analyses. It also details how a few careful considerations during the planning phase can lead to easier maintenance and greater reusability.

Spotfire provides a flexible and powerful framework for interactive visuals and, with some efficiency considerations, it can be used to handle complex analyses and provide key insights.

7. REFERENCES

[1] https://www.qualityforum.org/Projects/n-r/Patient-Reported_Outcomes/Patient-Reported_Outcomes.aspx
[2] https://www.cdisc.org/standards/foundational/adam
[3] https://deviq.com/don-t-repeat-yourself/
[4] https://qol.eortc.org/questionnaire/eortc-qlq-c30/

8. ACKNOWLEDGEMENTS

The author would like to thank Sohail Mulla and Sara Tadesse-Bell. They offered invaluable guidance on the needs of the PCOR team as it relates to the application. Their foresight and expertise helped immensely during the development process and led to a final product that is reusable and easily customizable.

Special thanks go to Jingyuan Chen, Aditi Qamra, Suwen Li, and Vibha Raghavan. They used the application for their own studies. This added to the application's generalizability and helped to emphasize how the tool was fit for reuse.

CONTACT INFORMATION

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

Matt Buckley
Hoffman-La Roche Limited
7070 Mississauga Road
Mississauga, Ontario, Canada L5N 5M8
Email: matt.buckley@roche.com