Paper OS06

Let's not get lost in translation: open-source adoption at Novartis.

Orla Doyle, Advanced Quantitative Sciences, Novartis, Dublin, Ireland.
Mark Baillie, Advanced Quantitative Sciences, Novartis, Basel, Switzerland.

ABSTRACT

Open-source technologies are gaining traction in clinical trial analysis and reporting. At Novartis, we are modernizing our approach through two key efforts: (i) the increased adoption of open-source tools ("the furniture") and (ii) the implementation of a next-generation Statistical Computing Environment (SCE) ("the house"). These efforts are driven not only by technological opportunity but also by the evolving scientific and regulatory landscape. For example, recent guidance such as the ICH E9(R1) addendum on Estimands [3] and the FDA's covariate adjustment [4] recommendations have introduced new analytical expectations that often require novel methods, many of which first emerge in open-source ecosystems or must be developed internally. This context has motivated a value-driven approach that avoids mere "code translation" and instead focuses on improving the quality, flexibility, and efficiency of analysis and reporting. We illustrate this approach through recent applications in phase III clinical trials piloting these changes, highlighting the journey through innovations such as Mosaic, Analysis Results Datasets (ARDs), and dynamic data exploration.

- MOSAIC, ARDS, OPEN SOURCE, RBMI, DMC

INTRODUCTION

The pharmaceutical industry is experiencing a notable shift in how clinical trial data is analysed, reported, and reviewed. At Novartis, we are modernizing our approach through two key efforts: the adoption of open-source tools ("the furniture") and the implementation of a next-generation Statistical Computing Environment ("the house"). Our motivation is rooted in the changing scientific and regulatory landscape. Guidance such as the ICH E9(R1) addendum on Estimands [3] and the FDA's covariate adjustment [4] recommendations have reshaped expectations for clinical trial reporting. These changes often necessitate the use of new statistical methods that are not readily available in commercial software. Open-source solutions, whether community-developed or internally implemented, offer the flexibility and innovation needed to meet these demands.

These initiatives are not simply technical upgrades; they represent an evolution in our processes and ways of working. By moving beyond traditional "code translation" exercises, we aim to unlock greater efficiency, reproducibility, and regulatory alignment in clinical reporting. Central to this journey are solutions such as *Mosaic*, a metadata-driven automation framework for standard outputs, and a Data Monitoring Committee (DMC) App, which enables dynamic, validated safety data exploration for data safety monitoring boards [5].

Together, these tools illustrate how open-source adoption can drive meaningful improvements, streamlining routine tasks while empowering teams with flexible, scalable solutions. In the sections that follow, we detail the design, implementation, and impact of these technologies, and reflect on how they support our broader vision for modern, value-driven clinical quantitative science.

ORIGIN STORY: PILOTING IDEAS IN PHASE III CLINICAL TRIALS TO SPARK THE SHIFT TO OPEN-SOURCE

A recent Phase III clinical study served as a key pilot for Novartis' broader adoption of open-source tools. With a complex endpoint structure, including composite measures and patient-reported outcomes, and a treatment policy estimand that addressed key intercurrent events, the study posed unique analytical challenges. Regulatory expectations around missing data handling had evolved, and traditional methods such as MMRM were no longer considered adequate by health authorities [3].

To meet these demands, an internal review recommended the use of the {rbmi} R package for reference-based multiple imputation [6], a method not available in commercial software. The {rbmi} package aligned with software engineering best practices, offered rapid bug fixes, and benefited from strong community support. Its open-source nature enabled timely implementation and collaboration across internal teams—including Biostatistics, Statistical Programming, methodology, and software engineering, as well as engagement with external developers from the open-source community.

This pilot not only demonstrated the feasibility of using open-source tools for regulatory submission but also showcased the power of analysis results datasets (ARDs). The pilot experience laid the foundation for broader adoption of open-source solutions across therapeutic areas including the development of Mosaic, reinforcing the strategic value of flexibility, transparency, and reproducibility in modern clinical data science.

TOWARDS AUTOMATING STANDARD OUTPUTS

In clinical trial reporting, the generation of tables, figures, and listings (TFLs) is often one of the most resource-intensive and time-sensitive activities. Traditionally, this process has relied on manual programming of outputs, which can introduce variability, require extensive validation, and consume significant programmer and reviewer effort.

Automating the production of these standard outputs offers a clear opportunity to improve consistency and accelerate timelines. When the automation is deterministic and metadata-driven, the value is amplified: metadata defines both the structure and content rules of each output, ensuring that results are reproducible, traceable, and compliant with regulatory expectations. This approach minimizes the need for bespoke programming, reduces human error, and allows clinical teams to focus on study-specific insights rather than repetitive data preparation.

At Novartis, the adoption of open-source technologies has enabled a shift toward metadata-driven automation in clinical trial reporting. Central to this transformation is the use of Analysis Results Datasets (ARDs), standardised, analysis-ready outputs generated through deterministic processing of ADaM datasets [1]. These ARDs are orchestrated via the CDISC Analysis Results Standards (ARS) model, ensuring regulatory alignment, traceability, and reproducibility across studies. By embedding ARDs within the Mosaic framework, that aim is to reduce the need for double programming of standard outputs, significantly refocusing effort elsewhere while enhancing consistency and compliance. The ARD approach exemplifies how open-source tools, when paired with industry standards, can streamline reporting workflows and support scalable, validated clinical data science.

MOSAIC

Mosaic is an automation framework designed to streamline the generation of standard clinical trial tables, figures, and listings (TFLs). Built on deterministic, metadata-driven principles and leveraging industry standards such as CDISC's Analysis Results Standards (ARS), mosaic ensures that outputs are consistent, reproducible, and fully traceable.

The Mosaic framework follows a validation process that de-risks the outputs created using this workflow, which may enable us to eliminate double programming for these outputs.

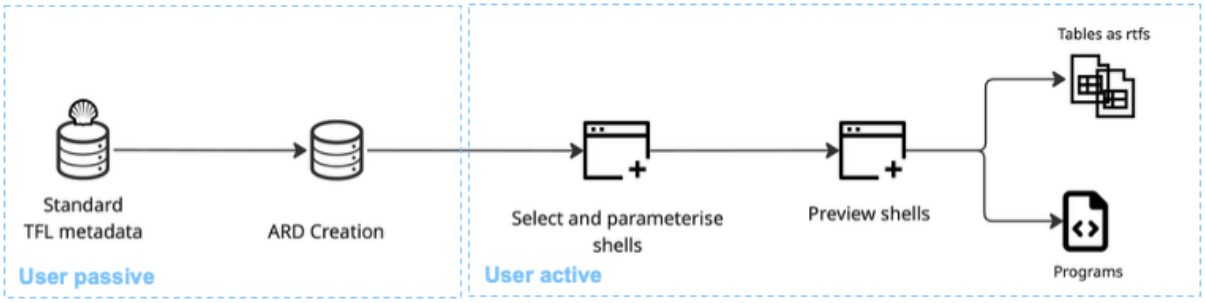

Step 1: Standard TFL Metadata

At the core of Mosaic is metadata that defines the structure and rendering instructions for standard TFL shells.

Standards compliance: Mosaic adopts the CDISC ARS model to ensure interoperability and regulatory alignment.

Coverage: The metadata represents different realisations of standard outputs. For e.g. an adverse events table could be realised with different AE filters (e.g. severe AEs vs all AEs). All options in scope of standard output are added to the metadata to ensure the downstream results will have complete coverage of the user's selections.

Implementation: A dedicated Python package encodes metadata and transforms it into the ARS data model.

Persistence: The ARS model is stored in a lightweight database, enabling portability, traceability, and ease of integration.

Step 2: Mosaic Engine (Computation Layer)

The Mosaic Engine transforms ADaM datasets into results datasets based on the metadata model.

Orchestration by ARS: The engine interprets the ARS model to determine the required operations.

Deterministic processing: All computations are metadata-driven, ensuring reproducibility and minimizing the risk of human error.

Output: The engine produces Analysis Results Datasets (ARDs), representing standardized, analysis-ready results that have full coverage of the options within the standard output.

Persistence: These ARDs are housed within the ARS and also stored as standalone parquet files.

Reproducibility: code is generated that will reproduce the results from the metadata.

Step 3: Rendering Engine (Output Layer)

Mosaic's rendering engine translates results into clinical reporting outputs.

Extraction: The engine selects relevant results from the ARS driven by the user selections.

Formatting: Outputs are rendered into tables according to predefined shells, preserving consistency across studies.

Reproducibility: code is generated that will reproduce each output starting from the ARS.

Extendibility: While initial focus is on tables, the same rendering approach can be applied to figures and listings.

User journey is enabled by the mosaic UI

During steps 1 and 2 mosaic operates in an automated manner where the user is passive. The user journey begins when they visit the mosaic UI. The UI facilitates user-driven realizations of our standard outputs.

Output Dashboard: mosaic provides the user with a dashboard of the outputs that they need to create for their reporting activity.

Configuration: mosaic exposes the different realizations of a standard output for user selection.

Preview: each output can be previewed in the app.

QC: the outputs can be QC-ed through a review workflow which any options are displayed to avoid unintended changes.

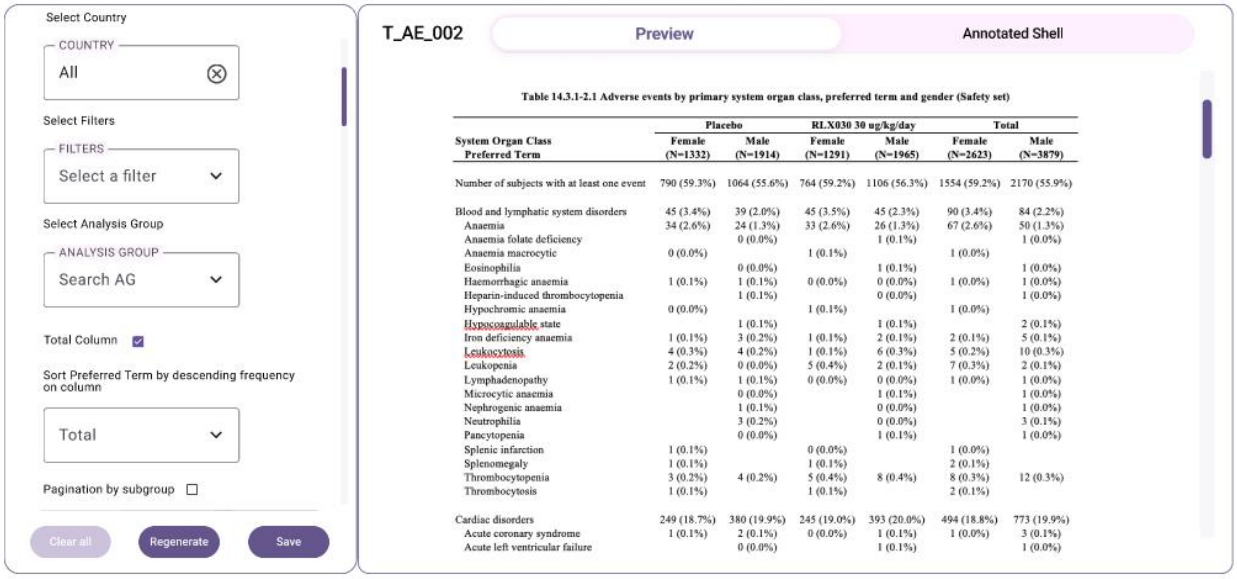

DYNAMIC DATA EXPLORATION

In clinical trials, timely and accurate safety monitoring is paramount. Data Monitoring Committees (DMCs) are entrusted with this responsibility, traditionally relying on static Closed Reports composed of pre-specified tables, figures, and listings. While these reports are essential for regulatory compliance and structured review, they often limit the ability to explore emerging safety signals [2].

To address this gap, we have developed a Dynamic Data Exploration App tailored for DMC use. This tool complements the Closed Report by offering an interactive environment where committee members can interrogate both patient-level and summary-level data. Built within a validated framework, the app ensures that flexibility does not come at the expense of reliability or compliance.

This section outlines the rationale, design, and implementation of the DMC App, highlighting how it supports enhanced safety oversight, scalable deployment, and study-specific customisation. It exemplifies how open-source innovation can empower stakeholders with tools that are both rigorous and responsive to the evolving needs of clinical trial governance.

DATA MONITORING COMMITTEE APPS

Data Monitoring Committees (DMCs) play a critical role in safeguarding participant safety and ensuring trial integrity. Traditionally, safety reviews have relied heavily on the DMC Closed Report, which presents static tables, figures, and listings. While essential, this format can limit the ability of committee members to probe deeper into emerging safety questions.

To address this, we have developed and validated a Dynamic Data Exploration (DMC) App designed to complement the Closed Report and provide a more flexible review environment. The application allows committee members to explore data interactively while ensuring alignment with the validated, pre-specified reporting framework.

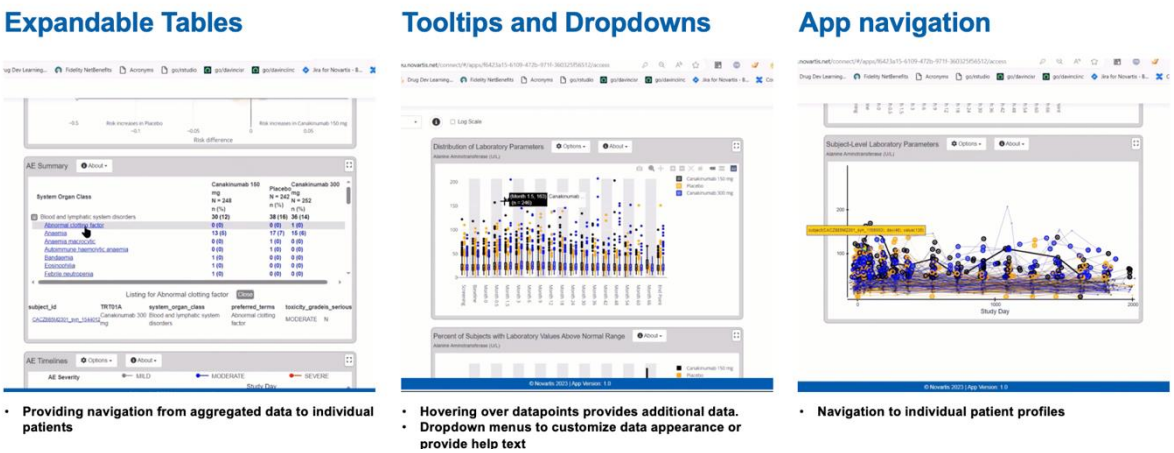

Benefits of the DMC App

Improved safety oversight: Interactive exploration of both patient-level and summary-level data enables timely and targeted investigation of potential safety concerns.

Validated and reliable: Developed under a controlled, validated process, providing sponsors and regulators with confidence in its outputs.

Efficient scalability: A streamlined deployment and validation model allows the app to be applied consistently across multiple studies and programs.

Tailored to study needs: Built-in flexibility supports study-specific customizations, ensuring relevance across diverse therapeutic areas.

Incremental Adoption Approach

Our journey toward implementing the DMC App has been intentionally incremental, ensuring that innovation was balanced with regulatory rigor and stakeholder confidence.

Pilot studies: Initial pilots were conducted across several studies to test the feasibility of dynamic data exploration within the DMC process. These pilots allowed us to refine workflows, gather committee feedback, and confirm that interactive outputs could be effectively integrated alongside the traditional Closed Report.

Validated toolkit development: Building on these learnings, we developed a validated toolkit that streamlined the deployment and validation of the DMC App. This approach drastically reduced the time and effort required to validate outputs, while fully maintaining quality and compliance standards.

Maturing to self-service: With the foundation in place, the process has continued to mature. Clinical trial teams are now adopting a self-service model, enabling study teams to configure and deploy the app without specialized coding expertise. This shift enhances scalability and empowers teams to focus on data interpretation rather than technical implementation.

This incremental approach has ensured not only the safe and reliable adoption of the DMC App but also its evolution into a sustainable, scalable solution that meets both regulatory expectations and the operational needs of sponsors.

CONCLUSION

The journey toward open-source adoption at Novartis is more than a technical transition, it is a strategic reimagining of how we analyse, report, and govern clinical trial data. Through initiatives ranging from phase III studies adopting R and ARDs **to the implementation of** Mosaic and the DMC App, we have demonstrated that metadata-driven automation and dynamic data exploration can deliver tangible improvements in efficiency, reproducibility, and regulatory compliance.

By embedding industry standards such as CDISC ARS into our workflows and embracing deterministic processing via ARDs, we have reduced manual effort, eliminated double programming, and empowered teams to focus on scientific insights. The incremental and validated rollout of the DMC App further exemplifies how innovation can be responsibly scaled across diverse studies.

These efforts reflect a broader cultural shift: from in-house legacy systems to open-source, value-driven solutions that support scalable, modern clinical data science. As we continue to build "the house" and furnish it with robust, interoperable tools, we remain committed to transforming not just our technology stack, but the very way we work.

REFERENCES

[1] Barros, J. M., Widmer, L. A., Baillie, M., & Wandel, S. (2022). Rethinking clinical study data: why we should respect analysis results as data. *Scientific Data, 9*, 686. https://doi.org/10.1038/s41597-022-01789-2

[2] Vandemeulebroecke, M., Baillie, M., Mirshahi, A., & Lesaffre, E. (2023). DMC reports in the 21st century: Towards better tools for decision-making. *Trials, 24*, 289. https://doi.org/10.1186/s13063-023-07290-4

[3] International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use. (2019). *Addendum on estimands and sensitivity analysis in clinical trials to the guideline on statistical principles for clinical trials E9(R1)*. https://database.ich.org/sites/default/files/E9-R1_Step4_Guideline_2019_1203.pdf

[4] U.S. Food and Drug Administration. (2023). *Adjusting for covariates in randomized clinical trials for drugs and biological products: Guidance for industry*. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/adjusting-covariates-randomized-clinical-trials-drugs-and-biological-products

[5] European Medicines Agency. (2005). *Guideline on data monitoring committees*. https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-data-monitoring-committees_en.pdf

[6] Gower-Page C, Noci A, Wolbers M (2022). "rbmi: A R package for standard and reference-based multiple imputation methods." *Journal of Open Source Software*, **7**(74), 4251. doi:10.21105/joss.04251, https://doi.org/10.21105/joss.04251.

CONTACT INFORMATION

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

Author Name: Orla Doyle
Company: Novartis AG
Email: orla.doyle@novartis.com

Author Name: Mark Baillie
Company: Novartis AG
Email: mark.baillie@novartis.com