PhUSE US Connect 2024
Paper OS03

NEST: What's Ahead?

Vincent Shen, Roche, Canada

Abstract

Over the past 7 years, NEST has evolved from a small proof of concept to an industry-recognized and well established project containing numerous R packages for streamlining analyses in clinical reporting and beyond. With large-scale adoption internally and ever-increasing engagement from external users, NEST is well positioned to further influence and revolutionize how we generate insights for clinical research in the pharmaceutical industry. In this paper, we will review our journey so far, share NEST's current capabilities and impact, and then highlight 3 prototypes that we are actively pursuing under the NEST umbrella: Multi-purpose Analysis Result Data, Regulatory Output Review Dashboard, and LLM-based NEST chatbot. Following that, we will discuss our long term vision on clinical insights generation and end the paper with a call on further and closer industry collaboration.

Glimpse of NEST's Journey

In 2015, a small group of people in the Biometrics (now Data Sciences) department at Roche started exploring the possibility of using R to conduct routine analyses for clinical trials. Various types of proof-of-concept (PoC) prototypes were set up and used for different purposes in clinical studies. Many of such prototypes were not scalable and with limited usability, but they started to trigger statistical programmers seeing the potential and benefits of R (compared to SAS), especially after R/Shiny had gained a great popularity in the R user group for interactive dashboard creation.

In 2017, the effort of building {teal} and {rtables} packages were officially kicked off by a team led by Adrian Waddell, and in less than a year, PoCs for both packages were launched and presented at the first ever R in Pharma conference. The potential capabilities offered by both packages surprised the audience from many pharmaceutical companies as well as regulatory agencies.

After the success at R in Pharma in 2018, biometrics leadership team at Roche decided to formally fund the PoCs, and this was when NEST has officially become a project. With improvement on infrastructure, a number of internally-developed R packages were deployed to an internal R server for exploratory analyses in clinical trials.

As the pharmaceutical industry is highly regulated, validation on computing environments and softwares used to generate clinical insights is a critical topic for consideration, and the potential shift to use open-source programming languages raised a lot of concerns from pharmaceutical companies. In 2020, after many rounds of discussions and clarifications, the first set of NEST packages were validated following Roche's internal process and were allowed to be used for regulatory reporting.

In 2021, with more companies from the industry also showing interest in developing and adopting R packages and open-source solutions for clinical insights generation, Roche made the decision to open source NEST codebase and invite industry peers to collaborate on the package development. The NEST codebase includes over 30 packages covering both regulatory and exploratory needs for clinical insights generation, and {rtables} made the first CRAN release in the same year.

Open-source gives NEST a lot more visibility in the industry and different components of the NEST project were highlighted in numerous internal and external events. In 2022, the project won the Roche PD Breakthrough Award, which was rarely given to a software engineering effort, and in the same year, NEST packages were used to support ad-hoc requests for a submission to FDA.

Over the years, Roche has made significant investments towards developing open-source solutions, not only for insights generation (by NEST), but also for creating SDTM and ADaM datasets. After many years of development, in 2023, these end-to-end open-source solutions have become the default toolkits to use for all new Roche trials, and at

the end of 2023, these open-source solutions have successfully delivered a readout for an R end-to-end Phase III trial, with the full R submission coming in 2024.

NEST now has over 1000 users, covering pharma, CRO and academia. The codebase has received contributions from over 200 developers across the world, and Roche has already partnered with 5 pharmaceutical companies to co-develop the NEST project. As a collection of open-sourced R packages, NEST aims to enable fast and efficient insights generation under clinical research settings, for both exploratory and regulatory purposes. To do that, we have open-sourced a catalog of 200+ tables, listings, graphs (TLGs) for clinical reporting, as well as over 70 built-in shiny modules for interactive analysis using {teal}. People can access the entire codebase on GitHub[1] and learn more about the entire NEST project on the NEST website[2].

NEST's Current Capabilities & Impact

Soon after NEST became a formally funded project, a main focus of the development work was to build template catalogs for common clinical trial analyses, so that users could easily adopt code templates and create TLGs. Many years of such effort led to two fully open-sourced, reproducible catalogs for TLGs in clinical insights generation, namely TLG-Catalog[3] and Biomarker Catalog[4]. Users learn the foundational NEST packages (e.g. {rtables}, {rlistings}, {tern}, {formatters}) through examples offered in both catalogs, and for many years, this approach was proven to be effective and also significantly lowered the bar for new R users to adopt these tools. Similarly, {teal}, a scalable modularized Shiny framework, offered a number of example apps in {teal.gallery}[5] so that it takes only a minimum effort to create a dashboard with dynamic filtering facility, code reproducibility, reporting engine, and many data summarization and visualizations.

Although individual templates and code examples offer great value for users who are new to NEST, there is a great need for a complete workflow to generate insights from ADaM datasets, with proper integration of metadata around study-specific information. Specifically, for a typical clinical trial, a list of planned analyses is created and agreed by all stakeholders, and programmers then develop programs to generate TLGs based on such a list. Many planned analyses are standard and even identical across studies, therefore, it is a common practice to manually copy over scripts from existing studies and then make study-specific adjustments. To reduce manual work and improve efficiency, {chevron} package was developed, which adopted a metadata-driven approach to automatically generate TLG scripts based on the TLGs on the list of planned analyses. The scripts are then integrated with another layer of metadata, and executed through Snakemake tools. This new workflow enables automated insights generation for all the standard analyses that require little or no customization, meaning that programmers could save significant amounts of time for delivering study-specific analyses. Such workflow is now in production use on the latest AWS-based computing environment at Roche, and has become the default solution for all new clinical studies.

A phase III trial adopted the workflow in late 2023, which allowed us to collect some metrics on how much efficiency gains have been achieved. Over 70% of the TLGs created for the study follow standards with no or little customization, and since scripts for those TLGs are automatically generated, programmers have saved at least 50% of time delivering such outputs. In addition, some TLGs were initially requested for making additional clinical decisions, and as these TLGs were mostly exploratory oriented and not part of the final clinical study report (CSR), a {teal} app was set up to answer ad-hoc questions from internal review. This was estimated to reduce at least 80% of turnaround time compared to generating static outputs for those requests.

What's Ahead?

With what we have already achieved today, NEST has been looking to continue our innovation journey and further improve efficiency on insights generation. We have built 3 prototypes and believe these efforts will redefine and revolutionize how clinical insights could be delivered.

Analysis-Result-Data (ARD)-Driven Insights Creation

[1] https://github.com/insightsengineering
[2] https://insightsengineering.github.io/nest/
[3] https://insightsengineering.github.io/tlg-catalog/
[4] https://insightsengineering.github.io/biomarker-catalog/
[5] https://insightsengineering.github.io/teal.gallery/demo.html

First, we are now building an open-source, multi-purpose analysis result data framework for harmonized insights generation in clinical study reporting. The prototype version of this framework (called {cards}) can be found on GitHub[6]. A little more than a year ago, 5 companies started an industry collaborative effort called {falcon}[7], with the aspiration to open-source a catalog of harmonized TLGs for clinical study reporting and simplifying the process of output review, re-use, and meta-analyses. The effort used FDA's proposed integrated guide for safety analyses as a reference, and 12 developers from these 5 companies together delivered 23 output templates since then. While crowd-sourcing templates creation is resource-efficient, industry collaborators also realized that it may not always be possible to standardize layout for clinical trial analysis, as different companies may have their own preferences on certain table/plot set-ups. Therefore, instead of aiming to harmonize layouts, {falcon} will adopt CDISC's recently proposed analysis result standard and use {cards} and {gtsummary} to build a multi-purpose ARD framework for clinical insights generation, including creation of an industry-version of TLG-Catalog. Once the framework is set up, insights would be first generated and stored in a universally agreed ARD format, and companies could then display results in various layouts and formats based on company-specific requirements.

Regulatory Report Dashboard

As we discussed earlier, {chevron} and the streamlined workflow built around it now enable automated script generation. It also integrates with study metadata, therefore offers great efficiency improvement in generating layouts and reports for standardized analyses in clinical studies. {teal}, on the other hand, was historically designed to be an exploratory framework to allow interactive data exploration, and its dynamic filtering feature greatly simplifies the process for ad-hoc analyses. There have been numerous discussions in the industry about delivering regulatory outputs using Shiny, but there are also concerns around this idea. One common feedback is that the dynamic filtering system provided in {teal} may be inappropriately used from a statistical perspective (e.g. a Kaplan-Meier plot for a subgroup with unknown age), and with little tracking on how a user filtered the data, this may lead to misinterpretation of a visualization. In another word, it is hard to connect the planned outputs for CSR (i.e. traditional static outputs) with all the possible ad-hoc insights delivered through a highly flexible dashboard.

In NEST, we have now started developing a framework for regulatory output delivery, review and exploration. This product will integrate {chevron} with {teal}, and combine the benefits from each of the packages. First, it displays results of every single planned analysis in a shiny dashboard. These outputs are directly generated by {chevron}, following the previously described automated workflow, and are easily identified from the built-in table of contents as well as the search functionality. It integrates the filter panel from {teal}, but only allows users to filter on pre-defined subgroups. These filters are defined through study-specific metadata and correspond to the statistical analysis plan. Since outputs directly come from {chevron}, users could select and download static outputs (together with code) if they prefer to keep a static copy of these results, but this product provides an integrated review and comment panel so that every stakeholder can log their comment directly within the dashboard by simply highlighting text on the outputs. This comment panel requires users to log in to use, and allow study teams to manage and trace feedback effectively. Compared to before, when stakeholders directly comment on PDF files or raise requests in emails, this product will significantly reduce the burden of stakeholder management, and enable scientists, statisticians, programmers to communicate much more efficiently.

On top of all these features, the script for creating such a dashboard is also automatically generated based on study metadata and those already auto-generated {chevron} scripts. As long as programmers create scripts (using {chevron}) for all the planned analyses, they can get this dashboard for free. We believe this is the gamechanger for our current process of insights delivery, review and exploration, and a prototype will be launched later this year.

LLM-based Chatbot

With all the hype around ChatGPT and LLM in the past year, NEST has been exploring different ideas around how to leverage these cutting-edge technologies. As Roche is transitioning from SAS to R for all clinical trial analyses and many programmers need to learn NEST and other open source R packages to deliver their daily tasks, we have built a prototype that serves as a personal assistant to help users learn and explore NEST packages. With GPT APIs, we successfully launched a chatbot internally, which was trained with curated code, templates, as well as comprehensive package documentation. The bot is able to create and modify NEST code and generate TLGs with natural language

[6] https://github.com/insightsengineering/cards
[7] https://pharmaverse.github.io/falcon/

prompts. It also provides extensive explanation on the code it modifies or creates, which could help new users a lot in their journey of transitioning from SAS-based tools to NEST. Currently, the chatbot only has knowledge on {rtables} and {tern}, but a plan has been in place to expand the knowledge base to other NEST products, as well as other R packages from the end-to-end process (e.g. {admiral}, {roak}).

Our Vision on Clinical Insights Generation and A Call for Collaboration

With the capabilities offered by NEST packages today, we could generate clinical insights faster than before. Open-sourcing all these packages also enable us to connect and collaborate with other industry peers to further accelerate the insights delivery process. Together with a modern computing platform, we believe NEST is set to deliver significant efficiency gains in clinical reporting. Putting things into metrics, we estimate that there will be a 90% time savings in creating standard TLGs and at least a 50% time savings in creating customized TLGs. However, we aim for things much bigger than this. With more adoption, process improvements, continuous innovation, collaboration, as well as new technologies, we envision to build an automated pipeline for multi-purpose insights delivery and redefine how the pharmaceutical industry creates and delivers clinical insights. The three prototypes we present in this paper all contribute to this long-term vision, and as what we have been doing for the past 3 years, we welcome our industry peers to join our journey and collaborate on these ideas. Ultimately, our aim is to develop free solutions for the entire clinical research community, making clinical reporting a much less resource intensive task. As a result, we could bring medicines much faster to patients.

Contact Information

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

Author Name: Vincent Shen
Company: Roche
Email: vincent.shen.vs1@roche.com