Paper TT02

R Validation: Getting Started

Christian S. Christensen, Lundbeck, Copenhagen, Denmark

ABSTRACT

'How do we start to use R on our clinical trials?' is a question currently relevant to many companies in the pharmaceutical industry. Along with this question comes considerations for validation, change management, infrastructure and generally handling of open-source languages within a clinical trial context.

The goal of this paper is not to provide an answer. Instead, it is to describe our ongoing journey at Lundbeck, a mid-sized pharma company, in the hope that our experiences and considerations may benefit you should you choose to pursue the answer yourself.

INTRODUCTION

Recent years have seen an increased interest in the use of R on clinical trial data. Alongside this development industry-wide initiatives are taking place to establish best practice and solutions which can lower the barrier of entry.

However, how can you approach this?

In this paper I will take it down a notch and share our progress and some of our challenges and considerations together with a set of resources that have been useful to us.

It is my hope that this will enable you to start out on your own journey more easily. Additionally, I hope it may contribute to an open conversation about the more practical challenges in getting started with R validation.

Note that the opinions expressed in this paper are those of the author and not necessarily Lundbeck.

GETTING STARTED

An initial challenge, prior to any formal project, is business support.

In Q4 of 2021 I had just joined Lundbeck as a Project Programmer. At this time SAS was used for all reporting purposes, but R was also used for some exploratory tasks – and preferred by some of our biostatisticians. My own professional experience was also primarily with SAS but with some excursions into Python and R.

At PHUSE EU Connect 2021 I attended an inspiring workshop on ADMIRAL, the ADaM in R Asset Library [1] which, aside from introducing the framework, also reinforced my opinion of the advantages in both the R language and open-source collaborations in pharma.

This, I decided, was too important not to take back to my organization. So, in December 2021, I shared the highlights and perspectives at our Biostatistics and Programming department meeting.

In beginning of 2022 a new Biometrics department goal was defined: 'Investigate the use of R for Clinical Trials'. From here it took some time but by Q3 2022 this led to the formation of our R Working Group with three members – a representative from biostatistics, data science and programming, each with 5-10% working time allocated to the project.

STARTING THE CONVERSATION

Initially we were not sure about the aim of the project, the stakeholders and what information was out there, so the first steps were exploration.

Reaching out internally to differing stakeholders to understand the interest, and limitations, in our organization, exploring online resources and regulatory guidance and getting a better understanding of both the business case and the scope of the project.

As a part of the process, we conducted a short internal survey to learn about the R capabilities and interest in R among our colleagues. The same survey also served as a basis for identification of 'R experts' who eventually would give input to the selection of packages.

Additionally, we benefitted from both external events such as PHUSE conferences and single day events and the R/Pharma conference as well as online resources such as:

- The 'R Adoption Series' by the R Consortium [2]
- The R Validation Hub website [3] and whitepaper [4].
- Transcelerate Modernization of Statistical Analytics (MSA) [5]
- The PHUSE R Package Validation Framework [6]

As for guidance regarding regulatory expectations the following were informative:

- ICH E9 [7], section 5.8.
- FDA Statistical Software Clarifying Statement [8]
- 21 CFR part 11 [9] & part 11 scope and application [10]
- R: Regulatory Compliance and Validation Issues [11]

Relatively early it became clear that we were looking to enable the use of R on clinical trials alongside SAS. Key motivators for the project were talent attraction and retention, use of fit-for-purpose tools (packages) in our clinical trials and a very real upcoming need for a primary analysis on a trial conducted in R.

In terms of change management, it should be mentioned that framing the conversation as 'enabling use of R *alongside SAS'* removed a lot of initial tension.

THE R ENVIRONMENT

Before progressing further, a suitable R environment had to be identified. It quickly became clear that we effectively had two options:

1. A cloud solution with Posit Workbench.
2. An upgrade of the Windows R Server connected to our existing Clinical Data Repository (CDR).

Benefits in the first were primarily the interface, scalability, and flexibility of the setup whereas the latter offered an easier-to-implement setup with integration directly into the clinical data repository.
It should be noted here that a cloud solution would not necessarily imply a clinical data repository in the cloud. Instead, our primary interest was in accessing data retained in the CDR for processing in the cloud – thereby lowering the added effort to GxP validate the setup.

In the end when the cases were presented to our management team the decision was to rely on the existing infrastructure. Deciding factors were the effort needed for the GxP validation and time due the upcoming need of a trial.

Alongside the discussions and investigation into the R environments we had settled on an approach to package handling. Perhaps contra-intuitively we decided on a larger set of packages than we expect to validate. The set was identified based on input from our R experts on future need, allowed for explorative analysis and may mean that less frequent updates to the configuration of packages would be needed.

Crucially, such a setup requires transparency towards users of what is validated.

To ensure reproducibility of the package installation an installation script was created in R. Initially we intended for an installation from source but due to issues ensuring the correct installation order for dependencies we opted for a mixed installation with most packages installed from binary.

UPGRADING THE R SERVER

While the installation script was being developed, we noted that the package versions were changing. Incorrectly we had assumed that installing from an older R version meant that package versions would be fixed.

The remedy to this was to install from a snapshot of CRAN hosted on MRAN, the Microsoft R Application Network.

At this stage we got increased involvement from our colleagues in the Biometrics systems department and progressed to initiate an upgrade of the SANDBOX environment.

This proved less straightforward than expected.

Up until this point installation and testing had been done by us and on our own laptops. Now it was by someone external to our company, in a different function, on a server, with different permissions and interfaces and no R experience. It also became clear that some settings, which had been defaulted in our local installations, had to be explicitly set.

In the end several iterations were required before we finally had the SANDBOX upgraded.

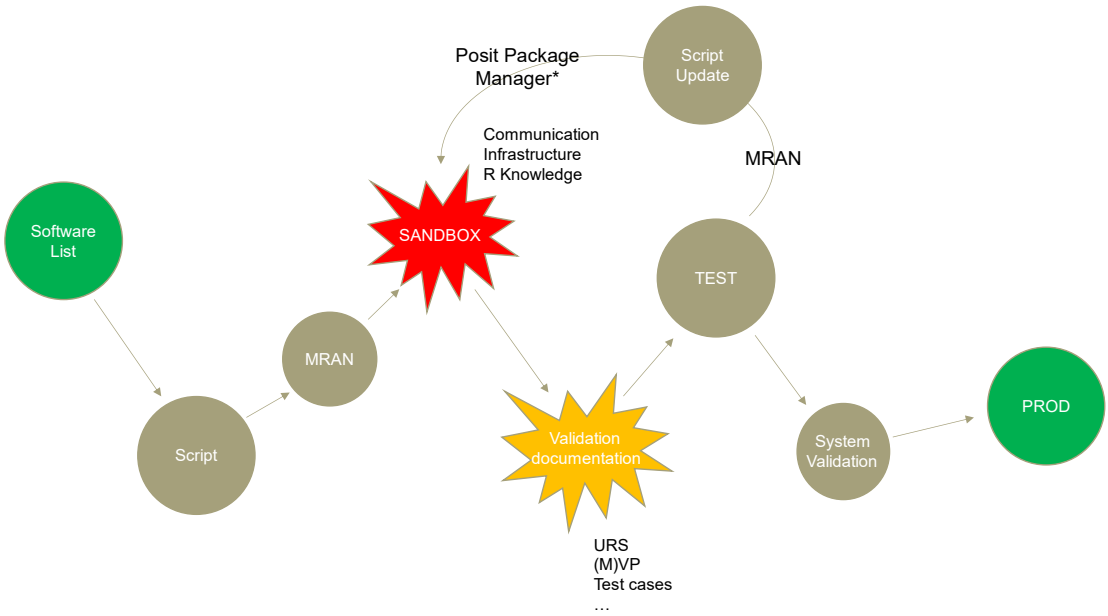

Figure 1: The R Server Upgrade.

After the SANDBOX upgrade, several formal requirements had to be completed before progressing to the TEST environment. This included formalizing User Requirement Specifications (URS), writing a Master Validation Plan, producing test cases and several rounds of review and approvals from relevant QA functions.

Two aspects made this substantially easier: The reliance on the existing validation of our CDR and the support from the Biometrics department responsible for our systems.

As we progressed to the installation in TEST however, the installation script was no longer executing as expected. MRAN had been taken down.

Luckily, we found a solution in the Posit Package Repository which led to an installation script update, a renewed installation in the SANDBOX environment, updates of the validation documentation and, finally, the intended installation in TEST.

It should be stressed here that the swap was not just a matter of replacing the package repositories.

First, despite using snapshots from the same date for the same R version, not all package versions on Posit Package Manager matched the packages previously available on MRAN. One may expect small deviations corresponding to packages changed within the day but in a few instances the package versions accessible on the Posit Package Manager were unlikely to have been available on the date of the snapshot.

Additionally, as could be expected from the description of the Posit Package Manager, a few additional packages had to be installed from source. A consequence of this was a renewed evaluation of the appropriate installation order.

To remove the external dependence of snapshots we strongly consider an installation from a local snapshot for any future updates.

After informal testing of the updated TEST environment, we proceeded to formal testing which progressed smoothly. Additionally, a window of opportunity for the update of our production environment has been identified – which, as I am writing this, is one week in the future.

PRINCIPLES FOR R PACKAGE VALIDATION

For the validation of R packages, we will be using a risk-based approach closely aligned with the descriptions in the R Validation Hub white paper and the Transcelerate MSA.

In this approach the focus will be on the directly used packages and not dependencies which are not directly called by the users. The basic idea here is that a validation of the directly used packages also tests the accuracy of any packages running back-end.
Moreover, this approach extends to R packages which are simply an interface to other code – such as packages relying on for instance Java or C++ code.

A basic validation flow is displayed in Figure 2. As can be seen the directly used packages go on to a risk assessment and, if the risk is medium or high, subsequent validation.

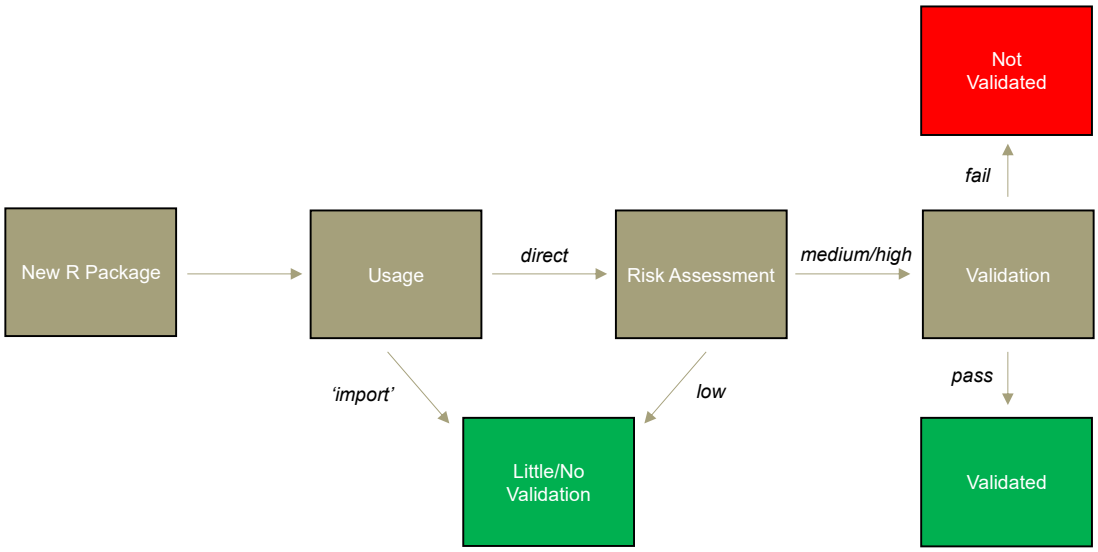

Figure 2: Validation Flow.

For the risk assessment we will rely on a grid adapted from the Transcelerate MSA. This grid, as displayed in Figure 3, uses 'confidence of accuracy' and 'impact of use' of a package together to assign an appropriate risk and validation level.

Confidence in Accuracy:

For our internally developed (SAS) programs we typically consider an assessment of complexity as a proxy for the risk of errors or mistakes occurring. However, open-source software comes with a history of use and therefore more information available. In addition to the complexity, it is therefore worthwhile to consider for instance community usage, level of documentation and testing and software maintenance when assessing confidence.

A very promising framework relevant in this context is the {riskmetric} package [12] [13].

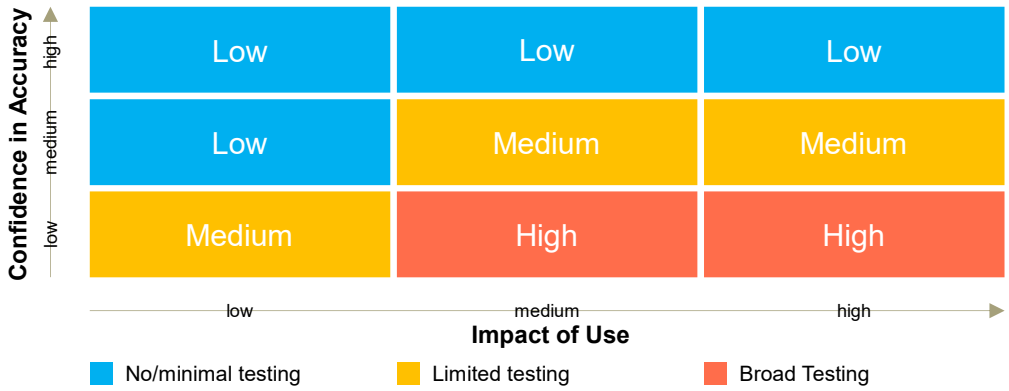

Figure 3: Risk assessment grid.

Impact of use:
In addition to the confidence, it is important to understand the possible consequence *if* a mistake or error occurs. Does it impact patient safety or a primary analysis (high impact)? Or would it merely affect exploratory analyses or data analyses for internal decision making (medium/low impact)?

In alignment with the discussion in [6] we consider base R and recommended packages to be highly trustworthy and do not consider any formal validation to be necessary.

PRACTICAL CONSIDERATIONS
While it will be too much to detail all the practical considerations that may go into package validation there are a few aspects which I will touch upon here.

One is the scope of the validation where one can consider a full validation of a package or only validating a subset of the functionality. This comes at a cost of transparency and a risk of improper use but lowers the burden of validating the package. An approach to mitigate the risk could be to create a wrapper code enabling only the functionality of interest from the original package.

Another aspect is lowering the effort needed to rerun a validation. Here it is worthwhile to create code for easy reuse in the case of system, R or package updates and automated tests, with expected, pre-specified outcomes, advantageous. To this end a framework such as {testthat} may be beneficial [14].

In some cases, however, it may not be so easy to see what agreement looks like. Consider the case of a numerical optimization with a random component. In this case results across languages, frameworks or packages are unlikely to agree completely as the result depend not only on the seed but also on the algorithm and the concrete implementations within the packages. In these cases, it is necessary to consider what an acceptable level of disagreement look like – and possibly also consider edge cases or asymptotic scenarios where solutions are simpler, and auto-correlation have less of an impact.

PROCESSES
As part of a final roll-out a standard operating procedure will need to be put in place to describe both the above detailed validation principles but also the necessary processes and documentation practice.

Here an important aspect is the anchoring of responsibility of the differing tasks.

In this respect our current intention is that it will be the R users who should bring forward their need for validation and drive the validation process forward. The R Working Group will then be represented by a system administrator who contributes to the risk assessment and approves the validation activities. Additionally, the system administrator archives the validation documentation and maintains an overview of the package validation status which is accessible to all R users.

FUTURE DEVELOPMENTS
An important next step, once we are in production, will be training. Training in proper use of the system but possibly also tutorials on how to do routine tasks in R and, for less straightforward tasks, reproduce results in SAS and R.

Additionally, we will have the first validations of R packages concluding and gain experience with the setup implemented; are there processes to be corrected or updates that need to happen for better use?

On other initiatives there has been interest in establishing a Lundbeck R user forum to facilitate the exchange of ideas and practices when working with R. This would be cross-functional and not restricted to Biometrics.

For the system setup we are strongly considering a local snapshot of packages in case of possible re-installation or future upgrades. This removes the dependency on online resources like Posit Package Manager and could be extended to include other software.
Likewise, in the case of future updates, it may be worthwhile to make the test cases more automated by reducing the manual steps involved.

Finally, alongside our project another platform is under development: a Lundbeck Cloud for Scientific Computing (LCSC). An interesting discussion will be how to align the two projects and in which way they may benefit from each other. To this end a data engineer, who is also engaged on the LCSC project, has recently joined the R Working Group.

CONCLUSION

As is hopefully clear from the paper this has been and is still a learning journey – and we have many interesting conversations ahead of us.

A few key takeaways are to not be afraid to pitch an idea to your organization, engage with users to identify requirements, ensure engagement from colleagues on the system-side to better discuss and prepare for system updates, not be discouraged by a bumpy ride – and finally to benefit from and contribute to the good discussions taking place currently.

I wish you all the best should you decide to get started yourself.

REFERENCES

[1] Multiple Contributors, "{admiral}: ADaM in R Asset Library," [Online]. Available: https://cran.r-project.org/web/packages/admiral/index.html.

[2] "R/Adoption Series," R Consortium, [Online]. Available: https://www.youtube.com/playlist?list=PL4IzsxWztPdkTYNmISc5Rd1kcBB6LQFgG.

[3] "R Validation Hub," [Online]. Available: https://www.pharmar.org/.

[4] R Validation Hub (subgroup), "A Risk-Based Approach for Assessing R Package Accuracy within a validated Infrastructure," 23 January 2020. [Online]. Available: https://www.pharmar.org/presentations/r_packages-white_paper.pdf.

[5] Transcelerate Biopharma Inc., "Modernization of Statistical Analytics," [Online]. Available: https://www.transceleratebiopharmainc.com/initiatives/modernization-statistical-analytics/.

[6] Phuse Working Group, "R Package Validation Framework," 15 DEC 2021. [Online]. Available: https://phuse.s3.eu-central-1.amazonaws.com/Deliverables/Data+Visualisation+%26+Open+Source+Technology/WP059.pdf.

[7] ICH, "E9: Statistical Principles for Clinical Trials," 05 FEB 1998. [Online]. Available: https://database.ich.org/sites/default/files/E9_Guideline.pdf.

[8] FDA, "Statistical Software Clarifying Statement," 06 MAY 2015. [Online]. Available: https://fda.report/media/109552/Statistical-Software-Clarifying-Statement-PDF.pdf.

[9] FDA, "21 CFR part 11," 20 MAR 1997. [Online]. Available: https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11.

[10] FDA, "Part 11, Electronic Records; Electronic Signatures - Scope and Application," SEP 2003. [Online]. Available: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/part-11-electronic-records-electronic-signatures-scope-and-application.

[11] The R Foundation for Statistical Computing, "R: Regulatory Compliance and Validation Issues," 18 OCT 2021. [Online]. Available: https://www.r-project.org/doc/R-FDA.pdf.

[12] Multiple Contributors, "{riskmetric} on CRAN," [Online]. Available: https://CRAN.R-project.org/package=riskmetric.

[13] E. Milliman and A. Clark, "{riskmetric} - R Validation Hub Mini-Series," [Online]. Available: https://www.youtube.com/playlist?list=PL4IzsxWztPdkuta1iqLmZckBdc1N4nZ9c.

[14] H. Wickham and and others, "{testthat}: Unit testing for R," [Online]. Available: https://cran.r-project.org/web/packages/testthat/index.html.

ACKNOWLEDGMENTS

I would like to thank the members of the Lundbeck R Working Group and key contributors to the project: Ingeborg H. Hansen, Mads M. Pedersen & Disha Singh for the continued support and effort to drive this project forward.

Additionally, I would like to acknowledge the effort of the many contributors across the industry who shares both considerations and solutions to the benefit of not only us but ultimately our patients.

A special thanks goes out to the R Validation Hub for the comprehensive overview which eases the entry to R validation and my colleagues at Novo Nordisk, Anders, Ari and Steffen, for sharing their experience and our discussions over the last two years.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Please contact me at:

Christian S. Christensen
H. Lundbeck A/S
Ottiliavej 9, 2500 Valby
crsi@lundbeck.com