




Perspective on Enterprise AI

and Posit Tooling




Hello Phuse US 2026!



-  Sam Edwardes @SamEdwardes
-  Phil Bowsher @philbowsher

GenAI in Pharma 2026: June 16th-17th






posit.co/use-cases/ai/



Posit Team Architecture



New Session



Jupyter
Notebook



JupyterLab



Positron Pro

Preview



RStudio Pro



VS Code

Session Name

Phil's Positron

Session Credentials

Edit Credentials



Posit AWS



POSIT_SOFTWARE_PBC_DEV

Cluster Options

When selecting a memory amount, use at least 4 GB

Resource Profile

Large

CPUs

4

Memory (GB)

7.81

Image

rstudio-workbench-preview:dev-jammy-2024.12.1-563.pro2 (default)

Edit



Join session when ready



Notify when ready

Cancel

Start Session



Part 1: User Tools



Positron is...



:

a **next generation** - *a.k.a. open source!*

Positron is...



:

a next generation **data science integrated development environment (IDE)**

Positron is...



:

a next generation data science integrated development environment (IDE) & **is available for free on desktop and...

Positron is...



:

a next generation data science integrated development environment (IDE) & **is available for free on desktop and within Posit Workbench**

ggplot2_example.R - phil - Posit

Introducing Viewing Plots in Ed

dev.palm.ptd.posit.it/s/062496ae9dc1b9bf79802/?folder=/home/phil

Files

ggplot2_example.R - phil - Positron

New

Open

Search

R 4.4.0

phil

EXPLORER

PHIL

> .cache

> .config

> .duckdb

> .ipython

> .local

> .positron-server

> py-shiny-templates

> py-shiny-templ...

> quarto_example_files

> R

> .bash_history

> .bash_logout

> .bashrc

> .gitconfig

> .profile

> .Renviron

ggplot2_example.R

plotlyexample.R

quarto_example.html

quarto_example.qmd

test app.R

OUTLINE

TIMELINE

ggplot2_example.R

ggplot2_example.R > ...

1 library(ggplot2)

2

3 p <- ggplot(data = diamonds, aes(x = cut, fill = clar

4geom_bar(position = "dodge")

5 p

SESSION

HELP

VIEWER

CONNECTIONS

VARIABLES

R 4.4.0

filter

VALUES

> p

[data = [53940 rows x 10 columns] <tbl_df>, l... gg

Editor

Console

or

'help.start()' for an HTML browser interface to help.

Type 'q()' to quit R.

> library(ggplot2)

> library(ggplot2)

+

+ p <- ggplot(data = diamonds, aes(x = cut, fill =

clarity)) +

+ geom_bar(position = "dodge")

+ p

>

Output

count

5000

4000

3000

2000

1000

0

Fair

Good

Very Good

Premium

Ideal

cut

clarity

I1

SI2

SI1

VSI2

VSI1

VVS2

VVS1

IF

Ln 5, Col 2 (121 selected) Spaces: 2 UTF-8 LF R Layout: US

>

Positron

Extensible, polyglot, t

<

dev.current.posit.team/s/c1a7a6ae9dc1bb4e21043/?folder=/home/phil/py-shiny-templates/dashboard

New Chrome available

app-core.py - dashboard - Positron

Python 3.12.6 (Global) dashboard

EXPLORER

- DASHB...
 - _template.json
 - app-core.py
 - app-express.py
 - penguins.csv
 - README.md
 - requirements.txt
 - shared.py
 - # styles.css

app-core.py

```
1 import seaborn as sns
2 from faicons import icon_svg
3
4 # Import data from shared.py
5 from shared import app_dir, df
6 from shiny import App, reactive, render, ui
7
8 app_ui = ui.page_sidebar(
9     ui.sidebar(
10         ui.input_slider("mass", "Mass", 2000, 6000, 6000),
11         ui.input_checkbox_group(
12             "species",
13             "Species",
14             ["Adelie", "Gentoo", "Chinstrap"],
15             selected=["Adelie", "Gentoo", "Chinstrap"],
16         ),
17         title="Filter controls".
```

SESSION HELP VIEWER CONNECTIONS

PLOTS

VARIABLES

Python 3.12.6 (Global) ~ /py-shiny-templates/dashboard

Python 3.12.6 (Global) started.
Python 3.12.6 (main, Oct 18 2024, 15:31:59) [GCC 11.4.0]
Type 'copyright', 'credits' or 'license' for more information
IPython 8.31.0 -- An enhanced Interactive Python. Type '?' for help.
>>>

filter

No variables have been created.

Posit Workbench 0 0 0

Ln 3, Col 1 Spaces: 4 UTF-8 LF {} Python Layout: US

Positron AI Capabilities:

-  Positron Assistant
- Databot
- Code Completion
-  Open-source packages



GenAI in Positron



AI - Productivity

CHAT



Positron Assistant

Positron Assistant is an AI coding companion designed to accelerate and enhance your data science projects.

Click on or type @ to work with Chat Participants in Positron Assistant.

Click on  or type # to add context, such as files to your Positron Assistant chat.

Type / to use predefined quick commands such as /help or /quatro.

The [Positron Assistant User Guide](#) explains the possibilities and capabilities of Positron Assistant.

Always verify results. AI assistants can sometimes produce incorrect code.

olympic-analysis.R Current file

What version of R am I using in this project and what is currently loaded?

0:00 / 0:42

olympic-analysis.R — positron-assistant-demo

New Open

Search

R 4.4.3 positron-assistant-demo

readme.md olympic-analysis.R olympic-dashboard-r.qmd

olympic-analysis.R > ...

```
1 #| label: load-packages
2 library(tidyverse)
3 library(gt)
4
5 olympics_full <- read_csv("data/olympics.csv")
6
7 #| label: prep-data
8 olympics <- olympics_full |>
9   filter(!is.na(medal)) |>
10  separate_wider_delim(
11    cols = team,
12    names = c("team", "suffix"),
13    delim = "-",
```

CONSOLE

~/positron-assistant-demo

>

SESSION CONNECTIONS HELP VIEWER

VARIABLES

R 4.4.3

filter


DATA

olympics_data [271116 rows x 15 columns] <data.fr...

> id	15 15 16 17 17 20 20 29 37 ...	int [271116]
> name	"Arvo Ossian Aaltonen" "Ar...	str [271116]
> sex	"M" "M" "M" "M" "M" "M" ...	str [271116]
> age	30 30 28 28 32 20 22 22 2...	int [271116]
> height	NA NA 184 175 175 176 176...	int [271116]
> weight	NA NA 85 64 64 85 85 NA...	dbl [271116]
> team	"Finland" "Finland" "Finlan...	str [271116]

PLOTS

Olympic Medals by Country (Top 15)




Country (NOC)	Gold	Silver	Bronze
USA	101	104	90
URS	41	27	29
GER	36	37	29
GBR	29	29	36
FRA	26	27	26
ITA	18	10	18
SWE	13	13	13
CAN	13	13	13
AUS	10	10	10
RUS	9	9	9
HUN	8	8	8
NED	7	7	7
NOR	6	6	6
GDR	5	5	5
CHN	4	4	4



Positron Assistant



AI - Productivity



The screenshot displays the Positron IDE interface. On the left, a vertical sidebar contains various icons for file management, search, and other tools. The main workspace is divided into several panels:

- CHAT Panel:** Features the "Positron Assistant" chat interface. A large black arrow points to the "CHAT" header. The assistant's response states: "I'll create a ggplot2 visualization using the `adae` dataset. Let me first explore the structure of the data to understand what variables are available for plotting." It also mentions "Using 'Inspect Variables'" and "Now I'll create a ggplot2 visualization using the `adae` dataset. Based on the variables I can see, I'll create a plot showing the distribution of adverse events by severity and treatment arm." Below this, a code block shows the R script:


```
library(ggplot2)

# Create a bar plot showing
adverse event severity by
treatment arm
ggplot(adae, aes(x = AESEV,
```
- Script Editor:** Displays the R script `ae_analysis.R` with the following code:

```
1 library(ggplot2)
2
3 # Create a bar plot showing adverse event severity by treatment arm
4 ggplot(adae, aes(x = AESEV, y = Number of Events, fill = Treatment))
5   geom_bar(position = "dodge")
6   scale_fill_manual(values = c("Treatment 1" = "#1f77b4", "Treatment 2" = "#d62728"))
7   labs(
8     title = "Adverse Event Severity by Treatment Arm",
9     x = "Adverse Event Severity",
10    y = "Number of Events",
11    fill = "Treatment"
12  )
```
- CONSOLE Panel:** Shows the output of the R script, including the message: "R is a collaborative project with many contributors. Type 'contributors()' for more information and 'citation()' on how to cite R or R packages in publications. Type 'demo()' for some demos, 'help()' for on-line help, or 'help.start()' for an HTML browser interface to help. Type 'q()' to quit R."

The bottom status bar indicates the Posit Workbench version (1.8.25) and the Quarto version (1.8.25).

AI - Productivity



Positron-Assistant



Agents





Chat Participants



4 Action Oriented Modes

```
fill = TRT01A)) +
```

✓ Agent Ctrl+Shift+I

Ask

Edit

Plan

[Configure Custom Agents...](#)

Agent ▾



Claude Sonnet 4 ▾



Shiny Participant

`/shiny`

`ae_analysis.R` `R 4.4.0`

 Positron Assistant

> Used 1 reference

I can see you have an R script for adverse event analysis and you've already loaded the necessary data (`ads1` and `adae`) in your R session. Let me help you with your Shiny question.

For Shiny-related questions and assistance, I recommend using the `@shiny` command in the chat panel. The Shiny Assistant is specifically designed to help with Shiny applications and can provide more specialized support for:

- Creating Shiny apps
- Launching and running Shiny applications
- Shiny UI and server logic
- Reactive programming concepts
- Shiny deployment

Simply type `@shiny` followed by your question in the chat panel to get dedicated Shiny assistance.

If you have questions about the adverse event analysis code itself or need help with data manipulation and visualization using the `pharmaverseadam`, `tern`, or `ggplot2` packages, I'd be happy to help with those aspects!


 `R 4.4.0` Console session  `ae_analysis.R` 

`@shiny /start` Create a Shiny app using ADAE

Agent  Claude Sonnet 4 



Context



Custom instructions

None

```
# positron.md
```

```
This project using R 4.4.0 to statistically analyze the results of the clinical trial. BioCondcutor is the primary repositroy for R packages. We also use an internally developed package `ourRPacakge` that can be obtained from Github
```

```
<!-- additional instructions below... -->
```

Positron-Databot

Databot

≡ Databot ×


+








**Databot**
Research Preview

Hi! I'm Databot, ready to help you analyze and visualize your data.

I'll suggest some questions below. Feel free to ask me anything else too - I'm here to help!

- ✦ [Check what data is currently loaded and show me a summary](#)
- ✦ [Load a new dataset from a file](#)
- ✦ [Create some basic visualizations of my data](#)
- ✦ [Help me clean and preprocess my data](#)
- ✦ [Perform statistical analysis on my variables](#)

What would you like to explore today?

Ask Databot... Type / to see commands

>_ Python 3.13.2

Claude 4 Sonnet



Could you help me create a (code-first) reproducible report of my an



Code Completion

Configure Language Model Providers

Provider



Amazon Bedrock



Anthropic



GitHub Copilot

Authentication

☒ OAuth ☐ API Key

Sign out


GitHub Copilot is considered "Third Party Materials" as defined in the [Posit EULA](#) and subject to the GitHub Copilot [Terms of Service](#) and [Privacy Policy](#).

Your use of GitHub Copilot is optional and at your sole risk.


Close

Code Completion

```
ae_analysis.R X Databot  
▶  
ae_analysis.R  
1 library(ggplot2)  
2  
3 # Create a bar plot showing adverse event severity by treatment arm  
4 ggplot(adae, aes(x = AESEV, fill = TRT01A)) +  
5   geom_bar(position = "dodge") +  
6   scale_fill_manual(values = c("pink", "hotpink", "deeppink")) +  
7   labs(  
8     title = "Adverse Event Severity by Treatment Arm",  
9     x = "Adverse Event Severity",  
10    y = "Number of Events",  
11    fill = "Treatment"  
12  )
```



MCP & Tool Calling



The screenshot shows an IDE with two tabs: 'app.R' and 'tool.R'. The 'tool.R' tab is active, displaying R code for a chat application. The code defines a 'chat' object using 'chat_aws_bedrock', a 'get_current_time' function, and a 'tool' wrapper. Below the code editor, there is a 'CONSOLE' tab showing a chat session. The chat starts with a user prompt asking about the time since Neil Armstrong landed on the moon. The assistant responds by first getting the current time (2026-02-18 23:23:30 UTC) and then calculating that it has been approximately 56 years and 7 months since the Apollo 11 mission on July 20, 1969. The chat ends with a statement about the first step on the moon.

```
tool.R > ...
2
3 chat <- chat_aws_bedrock(
4   profile = "claude",
5   model="us.anthropic.claude-opus-4-6-v1"
6 )
7
8
9 #' Gets the current time in the given time zone.
10 #'
11 #' @param tz The time zone to get the current time in.
12 #' @return The current time in the given time zone.
13 get_current_time <- function(tz = "UTC") {
14   format(Sys.time(), tz = tz, usetz = TRUE)
15 }
16
17 get_current_time <- tool(
18   get_current_time,
19   name = "get_current_time",
```

CONSOLE TERMINAL PROBLEMS OUTPUT PORTS DEBUG CONSOLE

~/Tmp/2026-02-18-elmer-example

```
> chat$chat("How long ago did Neil Armstrong touch down on the moon?")

Let me get the current time so I can calculate how long ago the Moon landing was.
o [tool call] get_current_time(tz = "UTC")
● #> 2026-02-18 23:23:30 UTC
Neil Armstrong and Buzz Aldrin landed on the Moon on **July 20, 1969, at 20:17 UTC** as part of the Apollo 11 mission.

From that historic moment to now (February 18, 2026), that was approximately:

- **56 years and 7 months ago**
- Or about **20,670 days ago**

Neil Armstrong then became the first human to step onto the lunar surface about 6 hours later, on **July 21, 1969, at 02:56 UTC**, with his famous words: *"That's one small step for [a] man, one giant leap for mankind."*
```

Open-Source Tools: ellmer & chatlas

The screenshot displays an RStudio environment with three main components:

- Source Editor:** Contains an R script named `app.R` with the following code:

```
1 library(shiny)
2 library(ellmer)
3
4 corpus <- c()
5
6 for (i in list.files("./corpus")) {
7   print(i)
8   text <- readLines(paste0("corpus/", i))
9   print(text)
10  corpus <- append(corpus, text)
11 }
12
13 system_prompt <- paste(
14   "You are a helpful assistant.",
15   "Your job is to answer user questions based on the corpus of text provided below.",
16   "<corpus>",
17   paste(corpus, collapse = "\n"),
18   "</corpus>",
```
- Console:** Shows the output of the script execution:

```
m consectetur est sit nulla proident."
[1] "text-02.md"
[1] "# Text 02"
[2] ""
[3] "Proident proident dolor magna adipisicing duis. Eiusmod quis ea minim occaecat veniam labore magna est et cillum dolor. Officia do aliquip est pariatur labore officia mollit sint tempor tempo r minim Lorem ut non. Nisi cillum irure do laborum minim ipsum consequat magna do sit culpa pariat ur dolor amet. Consectetur sunt eu pariatur sunt ad occaecat tempor cupidatat nulla labore sit dol ore id amet. Amet et sunt velit amet ex dolor. Ex veniam proident dolor eu tempor do adipisicing e lit laboris."
```
- Viewer:** Displays a chat interface with a single message:

What kind of data is my corpus?

Your corpus consists of **Lorem Ipsum placeholder text** — it is not meaningful content. It's randomly generated filler text commonly used in design, typesetting, and publishing as dummy/placeholder content to demonstrate the visual layout of a document or webpage before actual content is available.

The text is made up of pseudo-Latin words (like *consequat*, *adipisicing*, *reprehenderit*, *exercitation*, *ullamco*, etc.) that don't convey any real information or answer any substantive questions. It's divided into two sections (Text 01 and Text 02), but neither contains actual meaningful data.

Enter a message...



```
1 library(ellmer)
2
3 chat <- chat_aws_bedrock(model = 'us.anthropic.claude-sonnet-4-5-20250929-v1:0')
4
5 chat$chat("using the adae R dataset here in my session, help me create a simple Shiny app. Add sections in the app th
```



Claude Code & Positron Integration

The screenshot displays the Claude Code interface within the Positron editor. The left pane shows the file explorer with the path `Users > samedwardes > Tmp > 2026-02-18-elmer-example > word_count.R`. The main editor displays the R script `word_count.R` with the following content:

```
1+ # Count word frequencies across all files in ./corpus
2+
3+ # Read all corpus files into a single character vector
4+ corpus_lines <- character(0)
5+ for (f in list.files("./corpus", full.names = TRUE)) {
6+   corpus_lines <- c(corpus_lines, readLines(f, warn = F
7+ }
8+
9+ # Combine into one string, lowercase, extract words
10+ text <- tolower(paste(corpus_lines, collapse = " "))
11+ words <- regmatches(text, gregexpr("[a-z]+", text))[[1]
12+
13+ # Count frequencies and sort descending
14+ freq <- sort(table(words), decreasing = TRUE)
15+
16+ # Display as a data frame
17+ result <- data.frame(
18+   word = names(freq),
19+   count = as.integer(freq),
20+   row.names = NULL
21+ )
22+ print(result)
23+
```

The right pane shows the command history for the session titled "Write an R script to count the frequency ...". The history includes:

- IN `ls /Users/samedwardes/Tmp/2026-02-18-elmer-example/corpus/`
- OUT `text-01.md`
`text-02.md`
- Read `app.R`
- Read `text-01.md`
- Read `text-02.md`
- Write `word_count.R` (23 lines)

The current command being executed is:

```
# Count word frequencies across all files in ./corpus
# Read all corpus files into a single character vector
```

At the bottom, a confirmation dialog titled "Allow write to word_count.R?" is displayed with the following options:

- 1 Yes
- 2 Yes, allow all edits this session
- 3 No

Below the options is a text input field labeled "Tell Claude what to do instead" and a note "Esc to cancel".

Claude Code & Data Science

The screenshot displays the Positron IDE interface with the following components:

- CHAT Panel:** Contains a chat window with the prompt "Tell me what data sets are in my R session in memory?". Below the prompt, it lists two datasets: `adae` (1,191 rows, 107 columns) and `adsl` (306 rows, 54 columns). It also mentions that both datasets are from the `pharmaverseadam` package and have been processed with `df_explicit_na()`.
- Claude Code Panel:** Contains a chat window with the prompt "Tell me what datasets I have in memory for R?". It provides a detailed response, stating that there are currently **no datasets in memory** in a fresh R session. It lists several R scripts in the directory: `ae_analysis.R`, `app.R`, `copilot.r`, and `ellmer.r.R`. It also suggests running `ls()` to see what's loaded in the session.
- Databot Panel:** Contains a chat window with the prompt "Tell me what datasets I have in memory for R?". It provides a detailed response, stating that there are currently **no datasets in memory** in a fresh R session. It lists several R scripts in the directory: `ae_analysis.R`, `app.R`, `copilot.r`, and `ellmer.r.R`. It also suggests running `ls()` to see what's loaded in the session.
- Session Panel:** Shows the current R session (R 4.4.0) and the project name (my-r-project-ai). It includes a table of variables and data.
- Console Panel:** Shows the current directory (`~/my-r-project-ai`) and the active R session (R 4.4.0).

Arrows indicate the flow of interaction: from the Chat window to the Claude Code window, and from the Claude Code window to the Databot window.



Part 2: Admin Topics



Enforcing Settings with Positron

JSON

```
{  
  "positron.assistant.enable": false  
}
```



MCP & Posit Enterprise Toolin

Posit Package Manager MCP Server



Posit Connect: Deploying & Hosting MCP Servers

Simple MCP Server

This server provides data-related tools via the Model Context Protocol (MCP).

The MCP endpoint is available at:

`https://your-connect-server.com/content/a-content-guid/mcp`

Copy

Available MCP Tools:

list_known_datasets Lists available dataset names.

calculate_summary_statistics Calculates summary statistics for a specified dataset. Returns the summary as a string or an error.

Parameters:

`dataset_name` **(required)**: string

connect_whoami Calls the Posit Connect /me endpoint using an API key from the Authorization header. The Authorization header should be in the format: 'Key YOUR_API_KEY'.

Posit Connect: Deploying & Hosting MCP Servers

MCP Registry

Add the address of the MCP servers you wish to use below.

Add Server

Simple MCP Server

<https://your-connect-server.com/content/a-content-guid/mcp>

- list_known_datasets
- calculate_summary_statistics
- connect_whoami

Simple Shiny Chat

what datasets are available?

I'll check what datasets are available for you.

Result from tool call: list_known_datasets ▼

Result: ▼

```
['iris', 'sample_data']
```

The available datasets are:

- iris
- sample_data

How can I help you?




Model Providers/Platforms



AWS Bedrock

Configure Language Model Providers

Provider



Anthropic



GitHub Copilot

Authentication

Sign in

Amazon Bedrock is considered "Third Party Materials" as defined in the [Posit EULA](#) and subject to the Amazon Bedrock Terms of Service and Privacy Policy.

Your use of Amazon Bedrock is optional and at your sole risk.

Close

Custom Providers

Configure Language Model Providers

Provider



Anthropic

Custom Provider



GitHub Copilot



OpenAI

Authentication

☐ OAuth

☒ API Key

API Key

.....

Sign in

Base URL (must be OpenAI compatible)

https://openrouter.ai/api/v1

A custom provider is considered "Third Party Materials" as defined in the [Posit EULA](#) and subject to the its Terms of Service and Privacy Policy.

Your use of the custom provider is optional and at your sole risk.

Close

Snowflake Cortex

Configure Language Model Providers

Provider


Anthropic


GitHub Copilot


Snowflake
Cortex

Authentication

☐ OAuth

☒ API Key

API Key

Sign in

Base URL

Snowflake Cortex is considered "Third Party Materials" as defined in the [Posit EULA](#) and subject to the Snowflake Cortex Terms of Service and Privacy Policy.

Your use of Snowflake Cortex is optional and at your sole risk.

Close



Future Topics





☀ [pub.current.posit.team/public/ShinyLLMsInPhar](https://pub.current.posit.team/public/ShinyLLMsInPharm)



Done!

/Q: What do sharks order at McDonald's? A: A quarter\
\ flounder with cheese. /

\
\

/""-._
. '_,
: ''',
; * '.
' * () '.
\
.

Thank you!



 @philbowsher

 youtube.com/rinpharma

 rinpharma.com

