

The Shiny AI Stack

shinychat, querychat, and the LLM
plumbing beneath

Carson Sievert & Garrick Aden-Buie
Shiny engineers @ Posit



The stack (oversimplified)

Program w/ LLMs



Build chat UIs



Domain specific
'agents'





Quick chat app

```
ellmer::chat_posit() ▶  
shinychat::chat_app()
```

≡ claude-sonnet-5 (Anthropic)

Tell me a dad joke

Here's one for you:

Why don't scientists trust atoms?

Because they make up everything! 😄

Want another one?

+ Enter a message...



...or build your own chat app

```
library(shinychat)

ui ← page_chat("My chat app")

server ← function( ... ) {
  client ← ellmer::chat_posit()
  chat_server("chat", client)
}

shiny::shinyApp(ui, server)
```





```
library(shinychat)

ui ← page_chat("My chat app")

server ← function( ... ) {
  client ← ellmer::chat_posit()
  chat_server("chat", client)
}

shiny::shinyApp(ui, server)
```




Welcome to posit::conf(2026)! 🎉 It's a great afternoon to jump in — the Shiny AI Stack talk with Carson Sievert and Garrick Aden-Buie just kicked off in Ballroom A-F, and there's a Posit Connect session on role-based reports happening right now too. What are you hoping to get out of the conference — sharpening a skill, exploring AI, or just seeing what's new?

Interview me and help me build my agenda from scratch →

What should I see next after the Shiny AI Stack talk? →

Where and when are meals served today? →

+ Ask about sessions, workshops, or build your schedule... ⬆

🗨 Friendly ▾



- Searching the conf schedule >

📅 Searched the schedule: 2026-09-15; 15:20-16:20 14 matching items >

Right after The Shiny AI Stack ends at 3:20pm, stay put in Ballroom A-F: the same "No bad vibes: Building agents you can trust" track continues with two talks that pair naturally with it.

- **AI Agents for Data Science: Posit Assistant, Commons, and Beyond** (Joe Cheng), 3:20–3:40pm, Ballroom A-F. Picks up where the stack talk leaves off, covering Posit's agentic tools on top of that LLM plumbing.
- **AI Governance in Posit Team: getting better results from agentic helpers** (Elliot Murphy), 3:40–4:00pm, also in Ballroom A-F.

Then head to the keynote, **How to Be Clear When You Can't Be Certain** with Christine Zhang, 4:15–5:15pm, back in Ballroom A-F.

Want me to add the rest of that track (or just the next talk) to your agenda?

Add the full "No bad vibes" track to my agenda →

What else is happening Tuesday afternoon? →

+ Ask about sessions, workshops, or build your schedule...





• Searching the conf schedule

Find the Shiny AI Stack talk and its time slot ▾

Tool call

```
search_schedule(text = "Shiny AI Stack", `*_intent` = "Find the Shiny AI Stack talk and its time slot")
```

Tool result

```
[
  {
    "origin": "412a4001-b1bc-4903-99b0-3f7c1494f716",
    "chunk_id": 16,
    "text": "Title: The Shiny AI Stack: shinychat, querychat, and the LLM Plumbing Beneath\nKind: talk\nSession: No bad vibes: Building agents you can trust\nSpeakers: Carson Sievert, Garrick Aden-Buie\nTime: 2026-09-15 15:00:00 CDT to 2026-09-15 15:20:00 CDT\nFormat: HYBRID\nRoom: Ballroom A-F\nAbstract: LLM-powered chat is becoming a core pattern in modern Shiny apps, but\nbuilding a good chat experience (i.e., streaming, history, file\nattachments, editing, cancellation, etc) - is harder than it looks. In\nthis talk, we'll tour the Shiny AI stack: shinychat's new\ngetting-started API, which gives you a polished chat UI with all of\nthese niceties out of the box, and the lower-level plumbing beneath it\n(ellmer and chatlas) that handles the connection to the LLM. We'll also\nsee how these building blocks power domain-specific applications like\nquerychat, and why they're designed to be general enough for essentially\nany LLM-powered Shiny app. Examples in R, with Python equivalents\navailable for everything we'll cover."
  },
  {
    "origin": "412a4001-b1bc-4903-99b0-3f7c1494f716",
    "chunk_id": 17,
    "text": "Title: The Shiny AI Stack: shinychat, querychat, and the LLM Plumbing Beneath\nKind: talk\nSession: No bad vibes: Building agents you can trust\nSpeakers: Carson Sievert, Garrick Aden-Buie\nTime: 2026-09-15 15:00:00 CDT to 2026-09-15 15:20:00 CDT\nFormat: HYBRID\nRoom: Ballroom A-F\nAbstract: LLM-powered chat is becoming a core pattern in modern Shiny apps, but\nbuilding a good chat experience (i.e., streaming, history, file\nattachments, editing, cancellation, etc) - is harder than it looks. In\nthis talk, we'll tour the Shiny AI stack: shinychat's new\ngetting-started API, which gives you a polished chat UI with all of\nthese niceties out of the box, and the lower-level plumbing beneath it\n(ellmer and chatlas) that handles the connection to the LLM. We'll also\nsee how these building blocks power domain-specific applications like\nquerychat, and why they're designed to be general enough for essentially\nany LLM-powered Shiny app. Examples in R, with Python equivalents\navailable for everything we'll cover."
  }
]
```

+ Ask about sessions, workshops, or build your schedule...



 Searched the schedule: 2026-09-15; 15:20-16:20

Find what sessions start right after 3:20pm on Tuesday

		PM-4:00 PM	& I		Jana Sige, Thomas Lin Pedersen
Talk	AI Agents for Data Science: Posit Assistant, Commons, and Beyond	Tue Sep 15 3:20 PM-3:40 PM	Ballroom A-F	No bad vibes: Building agents you can trust	Joe Cheng
Talk	Quarto Wiki Chatbot	Tue Sep 15 3:20 PM-3:40 PM	Ballroom G & J	Organization-wide goodies	Aliza Samad, Sofia Lendahl
Talk	Bridging the Gap: Custom Catalogs to Connect End Users with Data Products	Tue Sep 15 3:20 PM-3:40 PM	Ballroom H & K	Dashboards	Malosree Maitra
Talk	A Grammar of Graphics for SQL	Tue Sep 15 3:20 PM-3:40 PM	Ballroom L & I	The Modern Data Science Stack... still SQL	Thomas Lin Pedersen
Talk	SQL Without Writing SQL: dbplyr for Serious Shiny Apps	Tue Sep 15 3:20 PM-3:40 PM	Ballroom L & I	The Modern Data Science Stack... still SQL	Brendan Broderick
Talk	AI Governance in Posit Team: getting better results from agentic helpers	Tue Sep 15 3:40 PM-4:00 PM	Ballroom A-F	No bad vibes: Building agents you can trust	Elliot Murphy
Talk	Positron for Everyone: Organization-Wide Posit Tooling on Amazon SageMaker	Tue Sep 15 3:40 PM-4:00 PM	Ballroom G & J	Organization-wide goodies	James Blair, Subodh Kumar
Talk	Template → Package → App: Meeting Users Where They're At	Tue Sep 15 3:40 PM-4:00 PM	Ballroom H & K	Dashboards	Jadey Ryan, Tanya Shapiro
Talk	ADBC: A Faster Foundation for Database Connectivity in R, Python, and Beyond	Tue Sep 15 3:40 PM-4:00 PM	Ballroom L & I	The Modern Data Science Stack... still SQL	Bryce Mecum, Ian Cook
Keynote	Keynote Session: How to Be Clear When You Can't Be Certain	Tue Sep 15 4:15 PM-5:15 PM	Ballroom A-F		Christine Zhang

Want me to add the rest of that track (or just the next talk) to your agenda?



Add the full "No bad vibes" track to my agenda →

What else is happening Tuesday afternoon? →

What talk is coming up next in this track?

🕒 Checked what's on now 3 items on now >

You're currently in **The Shiny AI Stack** (3:00–3:20pm, Ballroom A-F), which wraps in about 15 minutes. The next talk in the same "No bad vibes: Building agents you can trust" track is:

🖱️ **AI Agents for Data Science: Posit Assistant, Commons, and Beyond** — Joe Cheng, 3:20–3:40pm, Ballroom A-F.

No room change needed; after that comes Elliot Murphy's AI Governance talk at 3:40pm, then a break before the 4:15pm keynote.

Want me to add Joe Cheng's talk to your agenda?

+ Ask about sessions, workshops, or build your schedule... ⬆️

 New

🔍 Search...

TODAY

Recommendations A... 5:58 PM ...

Workshops Using AI 4:40 PM

Talks Covering New T... 4:40 PM

R vs Python Conferen... 4:38 PM

Want me to a

Add the full
agenda

🕒 Checked w

You're curre



(11:20)

- Lowering Barriers to Learning (Ballroom H & K): *Life after R4DS* (10:20), *Can I offer you a nice IDE?* (10:40)
- Quarto, Connected (Ballroom L & I): *Quarto + Accessibility + You* (10:20), *Slideware Is Dead* (10:40), *Quarto Hub* (11:00), *NASA Earthdata Cloud Cookbook* (11:20)
- Pretty Useful Data Viz (Ballroom G & J): *First Impressions Matter* (13:00), *Principles for Effective Interactive Visualizations* (13:20), *Do This Instead: 4 Alternatives to Common Charts* (13:40), *Reusing ggplot2 code* (14:00)
- From laptop to cluster (Ballroom H & K): *mirai + mori* (13:00), *targets and Hamilton* (13:20), *DuckLake in R* (13:40), *Tidymodels on Databricks* (14:00)
- Shining Stars (Ballroom L & I): *GeoGuessr in R* (13:20), *Open Source Routing Analysis at Scale* (14:00)
- Clinical Data Science (Ballroom G & J): *Environmental Stressors and Healthcare* (14:40), *Lowering Cesarean Rates with Quarto* (15:00), *ICU Dashboard at the VA* (15:20), *Ergonomic Informatics Solutions* (15:40)
- R Nerds (Ballroom H & K): *The Five Stages of R Optimization* (14:40), *Mocking Reality So Your Tests Can Run* (15:00), *When English Is Second, R Comes First* (15:20), *The Design of Everyday R Functions* (15:40)
- Developing the developer (Ballroom L & I): *It Worked Great Until I Left* (14:40),

Response cancelled



+ Ask about sessions, workshops, or build your schedule... 

What talks don't mention AI?



chedule x2



hose titles steer clear of AI, LLMs, agents, and prompting (note this is
tracks; a few abstracts may still touch on AI).

(Tiger Tang, 11:00–11:20 AM) maps second-language fluency to learning R, with a three-layer model that puts human judgment at the top.

- **Instant Feedback, Robust Evals: Agentic AI Tutoring in R and Python** (Michael Mullarkey, 11:20–11:40 AM) on giving learners immediate, well-evaluated AI feedback on their code.



Other learner-centered talks

- **Designing Reproducible Learning Ecosystems with R and Quarto** (Christian Martinez, Tuesday 2:00–2:20 PM, Ballroom H & K, "Teaching at the frontier" track) on growing learners into contributors through ownership and iteration.
- **Hello, robot: Introducing humans to LLM's** (Kristin Bott, Tuesday 11:00–11:20 AM, Ballroom A-F) on building AI skills by starting from learner curiosity and hands-on experimentation.

The Lowering Barriers to Learning track looks like the strongest fit for your interest, and it's only 40 minutes in one room. Want me to add it to your agenda?

Add the Lowering Barriers to Learning track to my agenda



Show me the full Teaching at the frontier track



What else is happening Wednesday morning?



+ Ask about sessions, workshops, or build your schedule...



What talks are about human-centered learning?

< 2 / 2



- Searching the conf schedule

Here are the talks that center on learners and how people learn:

Lowering Barriers to Learning track (Wednesday, Ballroom H & K)

- **From ABC to AI: Applying Language Acquisition to R and Technical Mastery** (Tiger Tang, 11:00–11:20 AM) maps second-language fluency to learning R, with a three-layer model that puts human judgment at the top.
- **Instant Feedback, Robust Evals: Agentic AI Tutoring in R and Python** (Michael Mullarkey, 11:20–11:40 AM) on giving learners immediate, well-evaluated AI feedback



What talks quack like a duck?

< 3 / 3 > ✎

- Searching the conf schedule >

One talk quacks, literally: **Quack to the Future: Building Traceable Data Lakehouses with DuckLake in R** by Travis Gerke. It's about building a lakehouse stack with DuckLake for transactional storage and column-level lineage tracking, all in dplyr-friendly workflows with versioned, auditable table writes. It's part of the "From laptop to cluster" track, Wednesday 1:40–2:00 PM in Ballroom H & K.

Quack to the Future: Building Tr...

If you enjoy that, you might also like the SQL-flavored **SQL as a first-class language in your data science IDE** in the "The Modern Data Science Stack... still SQL" track that Tuesday afternoon. SQL as a first-class language in ...

Show me the full "From laptop to cluster" track →

Add Quack to the Future to my agenda →

What else is happening Wednesday at 1:40 PM? →

+ Ask about sessions, workshops, or build your schedule...



- [Searching the conf schedule](#)

One talk quacks, literally: **Quack to the Future: Building Traceable Data Lakehouses with DuckLake in R** by Travis Gerke. It's about building a lakehouse stack with DuckLake for transactional storage and column-level lineage tracking, all in dplyr-friendly workflows with versioned, auditable table writes. It's part of the "From laptop to cluster" track, Wednesday 1:40–2:00 PM in Ballroom H & K.

[Quack to the Future: Building Tr...](#)

If you enjoy that, you might also like the SQL-flavored **SQL as a first-class language in your data science IDE** in the "The Modern Data Science Stack... still SQL" track that Tuesday afternoon. [SQL as a first-class language in ...](#)

Tue 2:40 PM–3:00 PM Ballroom L & I ·
Brian Lambert · Julia Silge

Show me the full "From
laptop to cluster" track

[Add to my agenda](#)

What else is happening
Wednesday at 1:40 PM?



What talks quack like a duck?

< 3 / 3 >

- Searching the conf schedule



One talk quacks, literally: **Quack to the Future: Building Traceable Data Lakehouses with DuckLake in R** by Travis Gerke. It's about building a lakehouse stack with DuckLake for transactional storage and column-level lineage tracking, all in dplyr-friendly workflows with versioned, auditable table writes. It's part of the "From laptop to cluster" track, Wednesday 1:40–2:00 PM in Ballroom H & K.

Quack to the Future: Building Tr...

If you
your
Tues

/now

What's happening right now?

/speaker

Show a speaker's card

/add

Find an item and add it to your agenda

/interview

Get interviewed about your interests to build your agenda

s a first-class language in
ick... still SQL" track that

Sho
lap

What else is happening
Wednesday at 1:40 PM? →

+ /◆

↑

DUCKLake in R

📅 2026-09-16 ⌚ 1:40 PM–2:00 PM 📍 Ballroom H & K

Modern data pipelines need more than clean outputs: they need a paper trail. This talk introduces two R packages to help you build a lakehouse stack using DuckLake for...

[Show more](#) ▾



Travis Gerke

Head of Applied Intelligence, PCCTC

[LinkedIn](#)

🔖 Added to the agenda

1 items in agenda >

Added **Quack to the Future: Building Traceable Data Lakehouses with DuckLake in R** to your agenda. It's Wednesday, September 16, 1:40–2:00 PM in Ballroom H & K, part of the "From laptop to cluster" track with Travis Gerke.

Quack to the Future: Building Tr...

Show me the other talks
in the "From laptop to
cluster" track



What else is happening
Wednesday at 1:40 PM?



Show my full agenda



+ Ask about sessions, workshops, or build your schedule...



💬 Friendly



My Agenda



Quack to the Future: Building Traceable Data Lakehouses with DuckLake in R



Wed, Sep 16 13:40–14:00 · Ballroom H & K

Quack to the Future: Building Traceable Data Lakehouses with DuckLake in R

📅 2026-09-16 ⌚ 1:40 PM–2:00 PM 📍 Ballroom H & K

Modern data pipelines need more than clean outputs: they need a paper trail. This talk introduces two R packages to help you build a lakehouse stack using DuckLak...

[Show more](#) ▾



Travis Gerke

Head of Applied Intelligence, PCCTC

[LinkedIn](#)

🔖 Added to the agenda

1 items in agenda >

Added **Quack to the Future: Building Traceable Data Lakehouses with DuckLake in R** to your agenda. It's Wednesday, September 16, 1:40–2:00 PM in Ballroom H & K, part of the "From laptop to cluster" track with Travis Gerke.

Quack to the Future: Building Tr...

Show me the other talks
in the "From laptop to
cluster" track



What else is happening
Wednesday at 1:40 PM?



Show my full agenda



+ Ask about sessions, workshops, or build your schedule...



💬 Friendly ▾

✎ New

🔍 Search...

TODAY

Quacking Talks 5:59 PM

Recommendations A... 5:58 PM ⋮

Workshops Using AI 4:40 PM

Talks Covering New T... 4:40 PM

R vs Python Conferen... 4:38 PM

Want me to add the rest of that track (or just the next talk) to your agenda?



Add the full "No bad vibes" track to my agenda →

What else is happening Tuesday afternoon? →

What talk is coming up next in this track?

🕒 Checked what's on now 3 items on now >

You're currently in **The Shiny AI Stack** (3:00–3:20pm, Ballroom A-F), which wraps in about 15 minutes. The next talk in the same "No bad vibes: Building agents you can trust" track is:

AI Agents for Data Science: Posit Assistant, Commons, and Beyond — Joe Cheng, 3:20–3:40pm, Ballroom A-F.

No room change needed; after that comes Elliot Murphy's AI Governance talk at 3:40pm, then a break before the 4:15pm keynote.

Want me to add Joe Cheng's talk to your agenda?

+ Ask about sessions, workshops, or build your schedule...



💬 Friendly ▾

📝 New

🔍 Search...

TODAY

Quacking Talks 5:59 PM

Recommendations A... 5:58 PM ...

Workshops Using AI 4:40 PM

Talks Covering New T... 4:40 PM

R vs Python Conferen... 4:38 PM



Where can I catch this talk?

📅 Searched the schedule: speaker: Emily Riederer


1 matching items >

KEYNOTE

Keynote Session: Context is the Control Plane

Context has always been critical to applied data work – long before artificial intelligence made it a buzzword. But why aren't we nearly as intentional when...

Show more ▾



Emily Riederer

LinkedIn

Emily Riederer's keynote, **Context is the Control Plane**, is Wednesday, September 16 at 9:00–10:00 AM in **Ballroom A-F** — the opening keynote of day two, so arrive a bit early for a good seat.

+ Ask about sessions, workshops, or build your schedule...



💬 Friendly



✎ New

🔍 Search...

TODAY

Quacking Talks5:59 PM

Recommendations A...5:58 PM⋮

Workshops Using AI4:40 PM

Talks Covering New T...4:40 PM

R vs Python Conferen...4:38 PM

of day two, so arrive a bit early for a good seat.

add both talks to my schedule

- ◆

Thought for 2s

>
- 🔖

Added to the agenda x2

>
- ◆

Thought for 6s

>

Done! Both talks are now in your My Agenda drawer:

- **AI Agents for Data Science: Posit Assistant, Commons, and Beyond** — Tue 3:20–3:40pm, Ballroom A-F
- **AI Governance in Posit Team** — Tue 3:40–4:00pm, Ballroom A-F

You'll also see the earlier items already on your schedule. Want anything else added, like the 4:15pm keynote?

Add Christine Zhang's 4:15pm keynote

→

Add Emily Riederer's Wednesday keynote too

→

+ Ask about sessions, workshops, or build your schedule...

⬆

💬 Friendly

⌵

My Agenda

✕

AI Agents for Data Science: Posit Assistant, Commons, and Beyond

Tue, Sep 15 15:20–15:40 · Ballroom A-F

✕

AI Governance in Posit Team: getting better results from agentic helpers

Tue, Sep 15 15:40–16:00 · Ballroom A-F

✕

Quack to the Future: Building Traceable Data Lakehouses with DuckLake in R

Wed, Sep 16 13:40–14:00 · Ballroom H & K

✕

🔍 Search...

R vs Python Conferen... 4:38 PM

add both talks to my schedule

◆ Thought for 6s

- **AI Agents for Data Science: Posit Assistant, Commons, and Beyond** — Tue 3:20–3:40pm, Ballroom A-F
- **AI Governance in Posit Team** — Tue 3:40–4:00pm, Ballroom A-F

Add Christine Zhang's 4:15pm keynote	Add Emily Riederer's Wednesday keynote too
--------------------------------------	--

+ Ask about sessions, workshops, or build your schedule... 

 Friendly 

X

Q Search & filter with SQL... / Total 30m 30m Now

10:13 0:15 10:20 10:25 10:30 10:35 10:40 10:43

⌵ Collapse ⌵ Clear Filter local data Column headers Pin parent spans

Today

10:14:31	posit-conf-chat	6 reactive_update	co.posit.r-package.shiny	3.73s
10:14:33	posit-conf-chat	32 session_start	co.posit.r-package.shiny	364ms
10:14:33	posit-conf-chat	6 reactive_update	co.posit.r-package.shiny	1.42s
10:14:53	posit-conf-chat	session_end	co.posit.r-package.shiny	9.64ms
10:17:40	posit-conf-chat	session_end	co.posit.r-package.shiny	6.06ms
10:18:16	posit-conf-chat	32 session_start	co.posit.r-package.shiny	1.44s
10:18:18	posit-conf-chat	9 reactive_update	co.posit.r-package.shiny	12.5s
10:18:20	posit-conf-chat	32 session_start	co.posit.r-package.shiny	341ms
10:18:21	posit-conf-chat	9 reactive_update	co.posit.r-package.shiny	13.0s
10:18:43	posit-conf-chat	19 reactive_update	co.posit.r-package.shiny	26.2ms
10:18:43	posit-conf-chat	15 observe event on_chat_user_input	co.posit.r-package.shiny	21.1ms
10:18:43	posit-conf-chat	14 ExtendedTask <anonymous>	co.posit.r-package.shiny	44.6s
10:18:43	posit-conf-chat	9 invoke_agent	co.posit.r-package.ellmer	44.4s
10:18:43	posit-conf-chat	1 chat z-ai/glm-5.3-flash	co.posit.r-package.ellmer	9.92s
10:18:53	posit-conf-chat	execute_tool skill	co.posit.r-package.ellmer	72.2ms
10:18:53	posit-conf-chat	1 chat z-ai/glm-5.3-flash	co.posit.r-package.ellmer	11.4s
10:19:05	posit-conf-chat	execute_tool skill	co.posit.r-package.ellmer	149ms
10:19:05	posit-conf-chat	execute_tool skill	co.posit.r-package.ellmer	120ms
10:19:05	posit-conf-chat	1 chat z-ai/glm-5.3-flash	co.posit.r-package.ellmer	22.3s
10:19:28	posit-conf-chat	3 reactive_update	co.posit.r-package.shiny	7.72ms
10:18:43	posit-conf-chat	observe on_stream_complete	co.posit.r-package.shiny	440μs
10:18:43	posit-conf-chat	1 observe <anonymous>	co.posit.r-package.shiny	1.39ms
10:18:53	posit-conf-chat	session_end	co.posit.r-package.shiny	5.59ms
10:32:55	posit-conf-chat	32 session_start	co.posit.r-package.shiny	1.55s
10:32:57	posit-conf-chat	9 reactive_update	co.posit.r-package.shiny	5.08s
10:33:22	posit-conf-chat	session_end	co.posit.r-package.shiny	6.37ms
10:35:20	posit-conf-chat	32 session_start	co.posit.r-package.shiny	332ms
10:35:21	posit-conf-chat	9 reactive_update	co.posit.r-package.shiny	4.60s
10:35:48	posit-conf-chat	15 reactive_update	co.posit.r-package.shiny	10.1s
10:37:00	posit-conf-chat	18 reactive_update	co.posit.r-package.shiny	11.1s

Automatically enabled

- History
- Editing
- Attachments
- Tools
- Thinking
- Citations
- Web search

```
library(shinychat)
```

```
chat_app()  
page_chat()  
chat_server()
```




A little customization

- Greetings
- Artifact Drawer
- Slash Commands
- Suggestions


```
library(shinychat)

chat_greeting()
chat_drawer()
$slash_command()
```






bit.ly/posit-conf-chat



opensource.posit.co/blog







Quick querychat app



```
library(palmerpenguins)


querychat_app(
  penguins,
  client = ellmer::chat_posit()
)
```

Quick querychat app



```
library(palmerpenguins)

querychat_app(
  penguins,
  client = ellmer::chat_pos
)
```





Welcome! I'm here to help you explore the penguin dataset—whether you're curious about species, island habitats, body measurements, or population patterns.

Explore the data

Show me the columns and a sample of the penguins data



What kinds of questions can I answer with this penguin dataset?



Analyze penguins

Compare average body mass across penguin species



Show how bill length and flipper length differ by species



+ Enter a message...



What kinds of questions can I answer with this penguin dataset?

Inspectable
SQL code

Fetch schemas: penguins 8 columns

Column	Type	Range / Values	Description	Constraints
species	TEXT	'Adelie', 'Chinstrap', 'Gentoo'		
island	TEXT	'Biscoe', 'Dream', 'Torgersen'		
bill_length_mm	FLOAT	32.1 to 59.6		
bill_depth_mm	FLOAT	13.1 to 21.5		
flipper_length_mm	INTEGER	172 to 231		
body_mass_g	INTEGER	2700 to 6300		
sex	TEXT	'female', 'male'		
year	INTEGER	2007 to 2009		

You can explore penguin **species, islands, physical measurements, sex, and year**. For example:

How many penguins are there by species?



What is the average body mass for each species?



Show a chart comparing bill length across species



+ Enter a message...





What is the average body mass for each species?

Query Data

Calculate average body mass for each penguin species

```
SELECT
  species,
  AVG(body_mass_g) AS avg_body_mass_g
FROM penguins
GROUP BY species
ORDER BY species
```

Result

	species	avg_body_mass_g
1	Adelie	3700.662
2	Chinstrap	3733.088
3	Gentoo	5076.016

Species

Average body mass

Adelie

3,701 g

Chinstrap


3,733 g

+ Enter a message...





Show body mass by species



ggsql code

+ Enter a message...





Save as PNG

Save as SVG

▼ Hide Query

↓ Save ▼

```
SELECT
  species,
  body_mass_g
FROM penguins
WHERE body_mass_g IS NOT NULL
VISUALISE species AS x, body_mass_g AS y
DRAW boxplot
LABEL x => 'Species', y => 'Body mass (g)'
```

ggsql code

+ Enter a message...




Generate a ridgeline plot of body mass. Place species on the y-axis and color each distribution by sex. Use a very transparent fill and colored stroke

◆ Thought for less than a second >

• Fetch schemas: penguins

8 columns >

📊 Body mass distribution by species and sex (ridgeline) ▾



> Show Query



↓ Save ▾

+ Enter a message...



So, you've generated some things,
now what?

Export insights with /handoff

≡ querychat with penguins

⚙️

📄

cohort.

3. **Adelie female, Biscoe 2009** — 20.7mm bill depth ($z = 3.3$), unusually deep for a female Adelie.

4. **Chinstrap female, Dream 2008** — only 2,700g ($z = 2.9$), the lightest bird in the dataset.

Pattern worth noting: Bill length is the most frequent "most extreme" measurement among the top outliers (5 of the top 12), and the two biggest outliers overall are both extreme bill lengths — consistent with bill length having the most within-group spread of the four measurements.

Most of these look like genuine biological variation rather than errors — the exceptions being the long-billed Chinstrap female and the lightest bird, which I'd verify if data quality matters for your use case.

Next steps

Filter the dashboard to show just these outlier penguins

Visualize bill length distributions by species and sex

Check if the extreme Chinstrap female overlaps Gentoo on other measurements

/handoff

Prepare a shareable handoff document or webapp using the current chat context.

+ /han

⬆️




Slash command

Select results, output format, etc



Prepare Handoff

Preserve important findings in a standalone report, dashboard, or script.

Results to include ⓘ



species	flipper...	body_...
Adelie	181	3750



Output format ⓘ

Quarto

Marimo

Shiny

Jupyter

... Other

Language ⓘ

Python

R

Generation notes ⓘ

⚡ PRE-FILLED BY AI

Build a Quarto dashboard organized into five sections that tell a complete story about the penguins dataset: (1) **The Dataset** — open with the dataset overview (query-0) alongside a data-quality card from the missing-records query (query-19), so readers know

⚡ Generate

View, revise, & download

querychat with penguins

- **Lightest:** A female Chinstrap at 2,700g (

Filter the dashboard to show only Gentoo penguins

Show me a chart comparing bill and depth with the Adelie species

Filter to rows with missing sex values so I can inspect them

Handoff

Quarto

A Quarto dashboard exploring the Palmer Penguin Paradox, with pages covering island/species overall vs within-species correlations, and the visualizations with DuckDB SQL query tables

+ Enter a message...

Handoff

title: "Palmer Penguins: Uncovering Simpson's Paradox"

format:

dashboard:

theme: cosmo

execute:

echo: false

warning: false

message: false

```\${r setup}

#| include: false

library(ggsql)

library(DBI)

library(duckdb)

library(dplyr)

library(knitr)

# Connect to an in-memory DuckDB instance and register the bundled CSV

# as the "penguins" table. Adjust the path below if penguins.csv is not

# located alongside this document.

con <- dbConnect(duckdb::duckdb())

penguins <- read.csv("penguins.csv")

dbWriteTable(con, "penguins", penguins)

````

Overview


Palmer Archipelago penguins were measured across three

Download/revise

Quick querychat app

```
library(palmerpenguins)

querychat_app(
  penguins,
  client = ellmer::chat_posit(
  )
)
```



Quick querychat app



```
library(palmerpenguins)

querychat_app(
  penguins,
  client = ellmer::chat_posit()
)
```

Add more context



```
library(palmerpenguins)

querychat_app(
  penguins,
  data_dict = "penguins.yaml",
  client = ellmer::chat_posit()
)
```


Run queries in your db



```
library(dplyr)

querychat_app(
  tbl(con, "penguins"),
  data_dict = "penguins.yaml",
  client = ellmer::chat_posit()
)
```

Pins also
supported!

Also polars,
ibis, etc.



Supports numerous tables



```
library(dplyr)

querychat_app(
  list(tbl(con, "t1"), tbl(con, "t2")),
  data_dict = "penguins.yaml",
  client = ellmer::chat_posit()
)
```

Custom app



```
library(querychat)
library(palmerpenguins)

qc ← QueryChat$new(penguins)

ui ← qc$page("Penguins Chat")

server ← function(input, output) {
  vals ← qc$server()
}

shiny::shinyApp(ui, server)
```


Reactively read data (filtered by LLM)



```
server ← function(input, output) {  
  vals ← qc$server()  
  
  output$plot ← renderPlot({  
    penguins ← vals$table("penguins")  
    filtered ← penguins$df()  
    hist(filtered$body_mass_g)  
  })  
}  
  
shiny::shinyApp(ui, server)
```

Curated insights

Hello! 🙌 Welcome to your Lending Club loan dashboard. I can help you filter, explore, analyze, and visualize your loan data with just a few clicks. Here are some ideas to get you started:

Filter

- Show small business loans
- Show loans that are most likely to default

Query

- Which loan purpose has the highest average interest rate?
- Do renters pay higher interest rates than homeowners?

Visualize

- Show how interest rates depend on loan grades
- Show how loan amounts have changed over time

Let me know what you'd like to do next!

+ Enter a message...



Total Loans
2,260,702




Average Interest Rate
13.09%




Total Loan Amount
\$34,016 MM


What Are People Borrowing For?




Loans by State



Better Credit Score, Lower Interest Rate



How Much Interest Do Borrowers Pay?



Querychat UI

Loan Data Explorer with QueryChat

Overview State Analysis Loan Details

Hello! 🖐️ Welcome to your Lending Club loan dashboard. I can help you filter, explore, analyze, and visualize your loan data with just a few clicks. Here are some ideas to get you started:

Filter

Show small business loans →

Show loans that are most likely to default →

? Query

Which loan purpose has the highest average interest rate? →

Do renters pay higher interest rates than homeowners? →

Visualize

Show how interest rates depend on loan grades →

Show how loan amounts have changed over time →

Let me know what I'd like to do next!

+ Enter a message...



Total Loans
2,260,702




Average Interest Rate
13.09%




Total Loan Amount
\$34,016 MM

What Are People Borrowing For?




Loans by State


Number of Loans



Better Credit Score, Lower Interest Rate



How Much Interest Do Borrowers Pay?



Hello! 🤖 Welcome to your Lending Club loan dashboard. I can help you filter, explore, analyze, and visualize your loan data with just a few clicks. Here are some ideas to get you started:

Filter

- Show small business loans →
- Show loans that are most likely to default →

Query


- Which loan purpose has the highest average interest rate? →
- Do renters pay higher interest rates than homeowners? →

Visualize


- Show how interest rates depend on loan grades →
- Show how loan amounts have changed over time →

Let me know what you'd like to do next!


+ Enter a message...




What Are People Borrowing For?




Loans by State



Better Credit Score, Lower Interest Rate



How Much Interest Do Borrowers Pay?



changed over time

Let me know what you'd like to do next!

Show small business loans

Fetch schemas: LOAN_DATA_FOCU... 51 colu...
Small Business Loans

```
SELECT * FROM LOAN_DATA_FOCUSED WHERE  
PURPOSE = 'small_business'
```

Apply Filter

The dashboard now shows only loans with
purpose = "small_business".

What is the average interest rate for small
business loans?

Show a chart of loan status breakdown for
small business loans

Compare small business loan amounts to
other purposes

Show me how default rate fluctuates
with FICO score.



Total Loans

24,689



Average Interest Rate


15.26%



Total Loan Amount

\$406 MM

What Are People Borrowing For?




Loans by State

Default Rate



Better Credit Score, Lower Interest Rate



How Much Interest Do Borrowers Pay?



Why restrict to SQL & ggsql?

Can guarantee read-only actions.

Strike a balance between safety & power

Want more capabilities? Add your own tools!

In summary

- Build your chat apps with shinychat + ellmer/chatlas
- Explore & augment curated insights with querychat



Thank you! Questions?

