

Twitter Thread by [Nick Singh | The Data Science Guy](#) ■



Nick Singh | The Data Science Guy ■
[@NickSinghTech](#)



A past coaching client of mine interviewed for an entry-level Data Science job in Los Angeles that paid \$110k per year.

Here's the 19 interview questions they were asked:

(how many of these could you answer?)

■■■■■■■■■■ ■

1. Explain the difference between population variance vs sample variance
2. Describe a p-value to a 13-year-old (8th grader)
3. What Does an R-Squared Value of 0.9 Mean?

■■■■■■■■■■ (Continued) ■

4. What are some metrics to evaluate how predictive your logistic regression model is?
5. Can you explain precision vs. recall? How does that relate to an ROC curve?

■■■ ■■

6. What is a database view?
7. What's the difference between a primary key & a foreign key?
8. Have you had to troubleshoot a slow query before? How did you do it?
9. When should you use a Sub-Query vs. CTE?

■■■ (Continued) ■■

10. Can you write a SQL query to find the 3-day rolling average?

(very similar to this Twitter SQL interview question on DataLemur)

<https://t.co/ssEQvcrKgo>

■■■■ (Continued) ■■

11. Can you write a SQL query to find the most frequently purchased pairs of items?

(very similar to the Walmart SQL Interview question on DataLemur)

<https://t.co/J8rwcjp8P5>

■■■■■■■ ■

12. Which Python libraries do you use?

13. What are some steps you take to clean dirty data?

14. In this dataframe, how would you remove all rows that met X condition?

■■■■■■■ (Continued) ■

15. Can you write code to find this value in a JSON object?

16. After loading the IRIS dataset via scikit-learn, can you build a simple regression model, and interpret the output?

■■■■■■■■■■■■■■■■■■■■ ■■■■

17. What are some ways you'd improve this sample Tableau dashboard, so a business stakeholder could get more value from it?

18. What metrics do you think are important to track for our business line? Do you see any potential flaws in those metrics?

■■■■■■■■■■■■■■■■■■■■ (Continued) ■■■■

19. Walk us through 2 of your past work projects.

What the most challenging technical and non-technical part of each project?

How many of these 19 questions could you answer?

Would you be able to land the \$110k Data Science job?

BTW if you enjoyed this thread:

1. Follow me [@NickSinghTech](#) for more data interview tips!
2. RT the tweet below to share this thread & challenge your followers! <https://t.co/GUkhNwMcRu>

A past coaching client of mine interviewed for an entry-level Data Science job in Los Angeles that paid \$110k per year.

Here's the 19 interview questions they were asked:

(how many of these could you answer?)

— Nick Singh | The Data Science Guy \U0001f4d5 (@NickSinghTech) [September 30, 2022](#)

Looking for resources to study for #DataScience & #MachineLearning interviews?

Follow this 12-Week Roadmap ■■

<https://t.co/NLJ4oU6NbL>

8 must-read Product Data Science books to level up & prep for your next Data Science or Product Data Analyst job interview:

(along with a guided 12 week study plan \U0001f447\U0001f447) pic.twitter.com/3cP2T7T911

— Nick Singh | The Data Science Guy \U0001f4d5 (@NickSinghTech) [May 2, 2022](#)

Also, don't forget, that #SQL is a KEY part of the Data Science and Data Analyst interview process.

Here's the 30 Day Plan to go from SQL Zero to SQL HERO!

■■

<https://t.co/7uRtIH1Y5F>

Learn SQL in 30 days for FREE with these 5 amazing resources:

(be sure to bookmark this 30-Day SQL Roadmap)

A thread\U0001f9f5\U0001f447

— Nick Singh | The Data Science Guy \U0001f4d5 (@NickSinghTech) [September 15, 2022](#)