VISCOSITY
PROFESSIONAL SERVICES

Crafting the Ultimate Outdoor Adventure Forecast Using MySQL HeatWave

MySQL and HeatWave Summit 2024 | May 1, 2024
Oracle Conference Center, Redwood City, California USA

MySQL HeatWave

ORACLE

Craig Shallahamer

Oracle ACE Director
Applied AI Scientist | OraPub Founder

@orapub
linkedin.com/in/craig-shallahamer-571a94a/
craig.shallahamer@viscosityna.com

Craig Shallahamer

Oracle ACE Director
Applied AI Scientist | OraPub Founder

@orapub
linkedin.com/in/craig-shallahamer-571a94a/
craig.shallahamer@viscosityna.com

Having worked with **Oracle technology since 1989**, Craig Shallahamer is a leader in the fields of machine learning, artificial intelligence and Oracle database performance tuning.

Craig has extensive experience in **constructing and teaching predictive modeling** methods, notably developing a Reinforcement Machine Learning bot in 1990. He has launched several **specialized generative AI conversational assistants**, each with distinct personalities and capabilities. As an Applied AI Scientist at Viscosity and the founder of OraPub, Craig is also the **author** of two acclaimed books: "Oracle Performance Firefighting" and "Forecasting Oracle Performance".

He has received recognition as an **Oracle ACE Director** for his valuable contributions to the Oracle community through his technical expertise and leadership. Craig is an active participant in Oracle user groups, **frequently presenting** at conferences and serving as a **board member and volunteer**.

1990

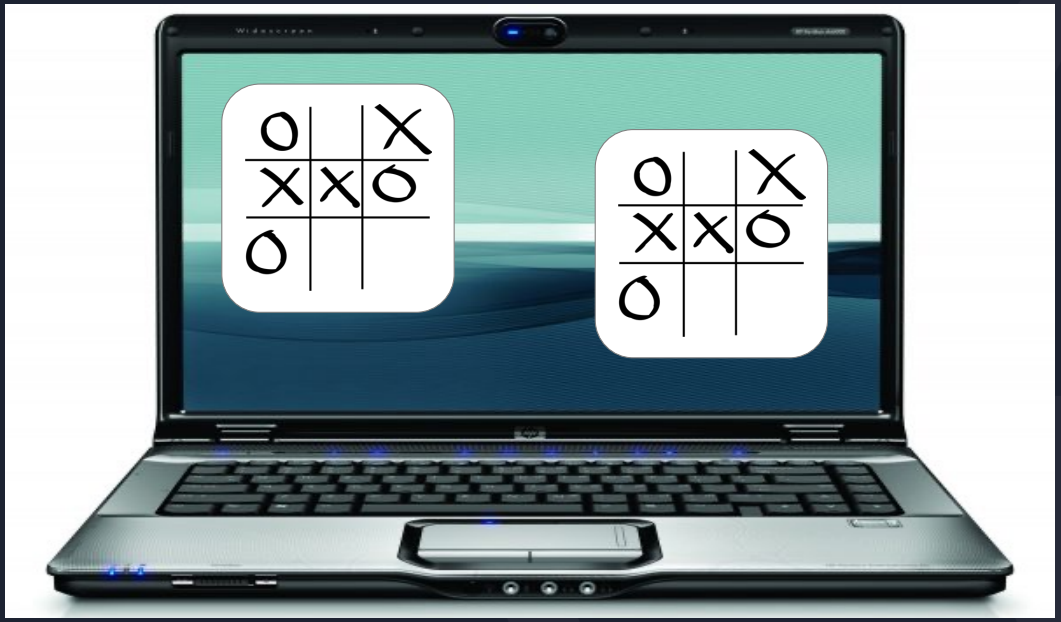

I created a program where I could play the game tic-tac-toe with a computer, but also enabled the computer to play itself… and learning through that process.

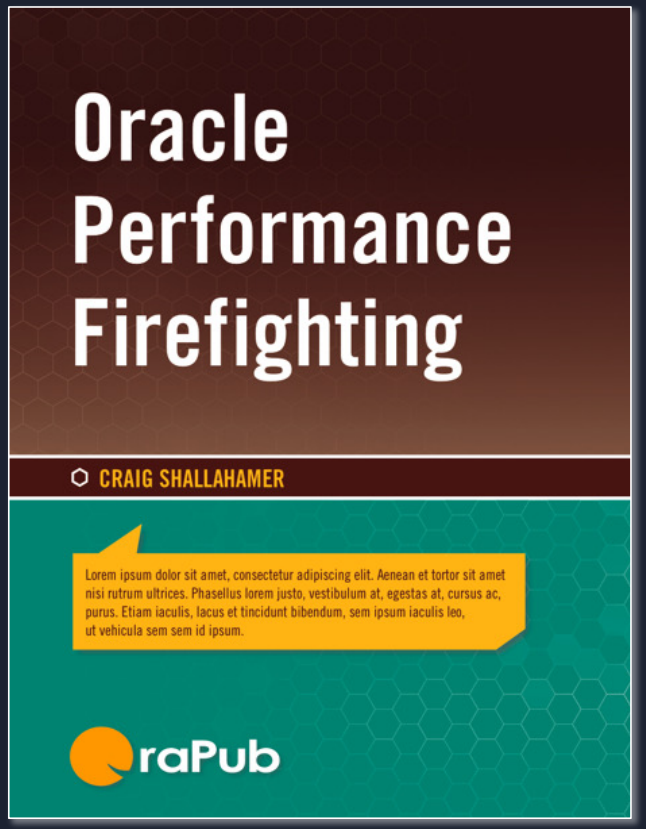
Oracle
Performance
Firefighting
CRAIG SHALLAHAMER
Lorem ipsum dolor sit amet, consectetur adipiscing elit. Aenean et tortor sit amet nisi rutrum ultrices. Phasellus lorem justo, vestibulum at, egestas at, cursus ac, purus. Etiam iaculis, lacus et tincidunt bibendum, sem ipsum iaculis leo, ut vehicula sem sem id ipsum.
raPub


THE EXPERT'S VOICE® IN ORACLE
Forecasting
Oracle
Performanc
Craig Shallahamer
Apress®
THE EXPERT'S VOICE® IN ORACLE
Forecasting
Oracle
Performance
Craig Shallahamer
Apress

MACHINE LEARNING 2

Performance Prediction

ORAPUB

MACHINE LEARNING 1

Anomaly Detection

ORAPUB

From a more business perspective...

- Internet marketing, focusing on identifying "hot leads" for very specific promotions.
- IT Operations when many Oracle systems must be monitored and Operations possibly alerted; both anomaly detection and performance prediction.
- Domain specific generative Q&A Bots and Conversational Assistants With Personality in multiple business domains. For example, Oracle Performance Tuning, Biblical Texts, Real estate investing platform.

Shane

Shane AI: The Oracle Tuning AI Assistant

Leveraging **Oracle Cloud Infrastructure**, AI Shane is a **conversational AI** assistant with a vast curated knowledge base of **all things Oracle performance**. Shane also has a very distinct stereotypic **surfer persona**. His objective is to help anyone **improve** their **Oracle system performance** optimization skills.

AI Shane leverages the **Oracle Generative AI** services within **Oracle Cloud Infrastructure** (OCI).

This includes **LLMs, vector** and traditional capabilities within Oracle Databases.

Specifically, AI Shane can use the Oracle Vector Store, Oracle GenAI Services, Oracle Autonomous DB, Oracle MySQL, Oracle MySQL Heatwave and OCI Compute.

ORAPUB
POWERED BY VISCOSITY

VISCOSITY

ORAPUB
POWERED BY VISCOSITY

VISCOSITY

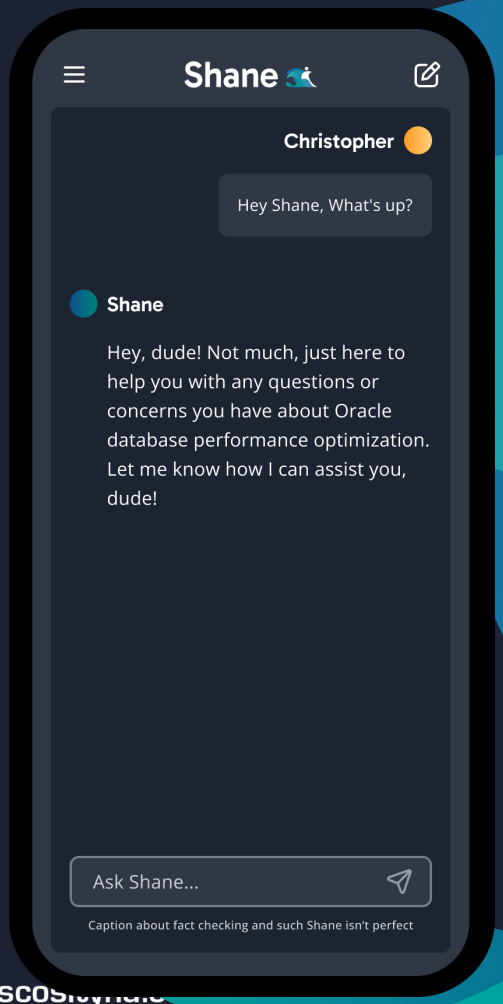

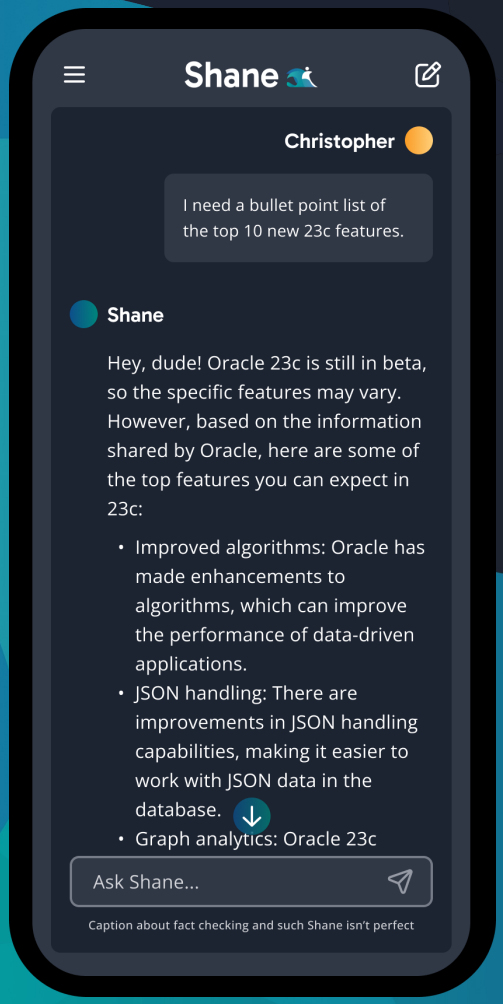

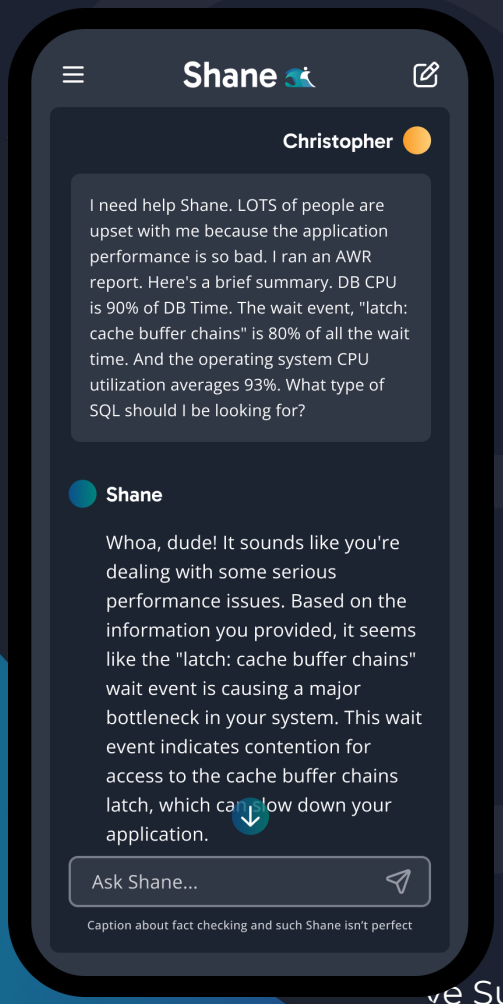

viscosity

ve Summit

Viscosity Pillars and Delivery Models

DATA

Oracle & SQL Server Postgres

Performance Tuning

Data Replication

Data Warehousing Analytics

Data Integration

ERP Blue Prints

Database Upgrades

APPS

APEX

EBS

Web/Mobile Apps

.Net and C#

E-Business Suite

SAAS/PAAS

CLOUD

Azure Gold Partner

Cloud Migrations

Engineered Systems

Oracle Cloud Partner

Google Partner

AWS Partner Hybrid Cloud

Workshops

Assessments

Proof of Concepts

Training

Turnkey Projects

Managed Services

ACE Program

Viscosity's Oracle ACEs

The Oracle ACE Program

The Oracle ACE Program recognizes and rewards individuals for their contributions to the Oracle community.

Charles Kim
CEO | Co-Founder

@racdba
ACE Director

Rich Niemiec
Chief Innovation Officer

@richniemiec
ACE Director

Craig Shallahamer
Applied AI Scientist

@orapub
ACE Director

Sean Scott
Principal Consultant

@oraclesean
ACE Director

Gary Gordhamer
Principal Consultant

@ggordham
ACE Pro

viscosityna.com

@ViscosityNA

MySQL and HeatWave Summit

We've written over 25 technical books!

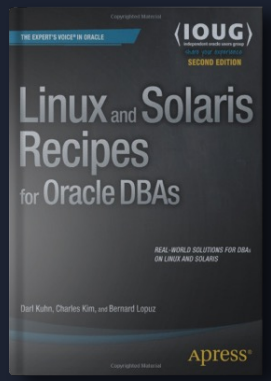

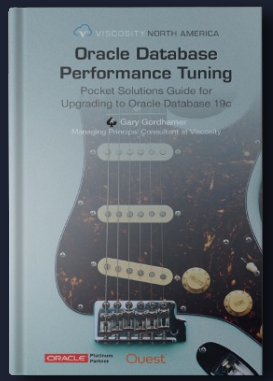

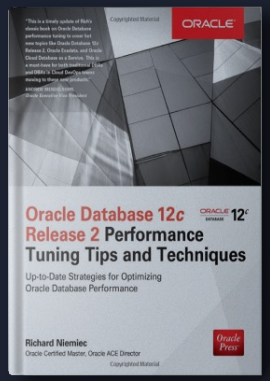

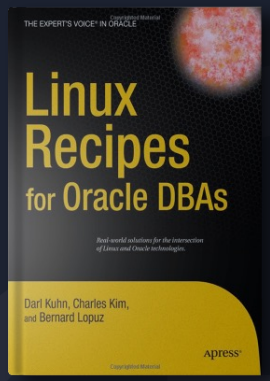

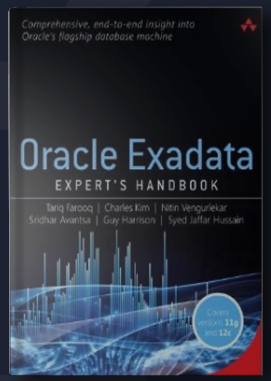

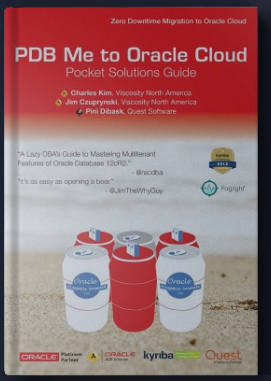

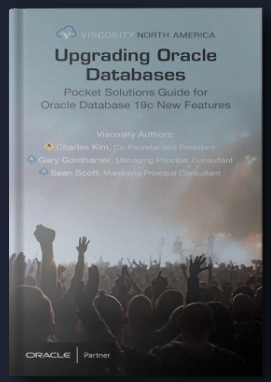

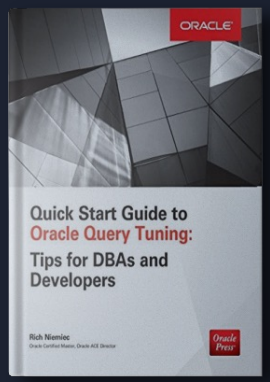

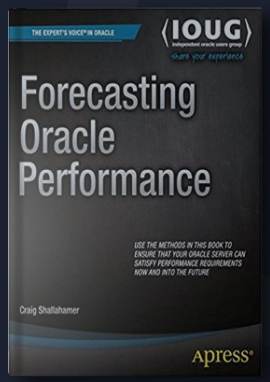

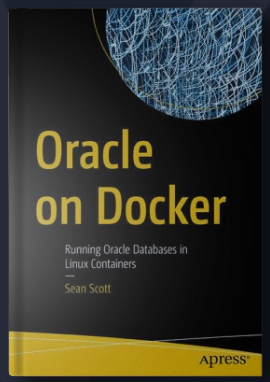


READY?
viscosityna.com
@ViscosityNA
MySQL and HeatWave Summit


THE PLAN
• Situation – the pain
• Flow
• Architecture
• MySQL & HeatWave calls
• Surprises
viscosityna.com
@ViscosityNA
MySQL and HeatWave Summit

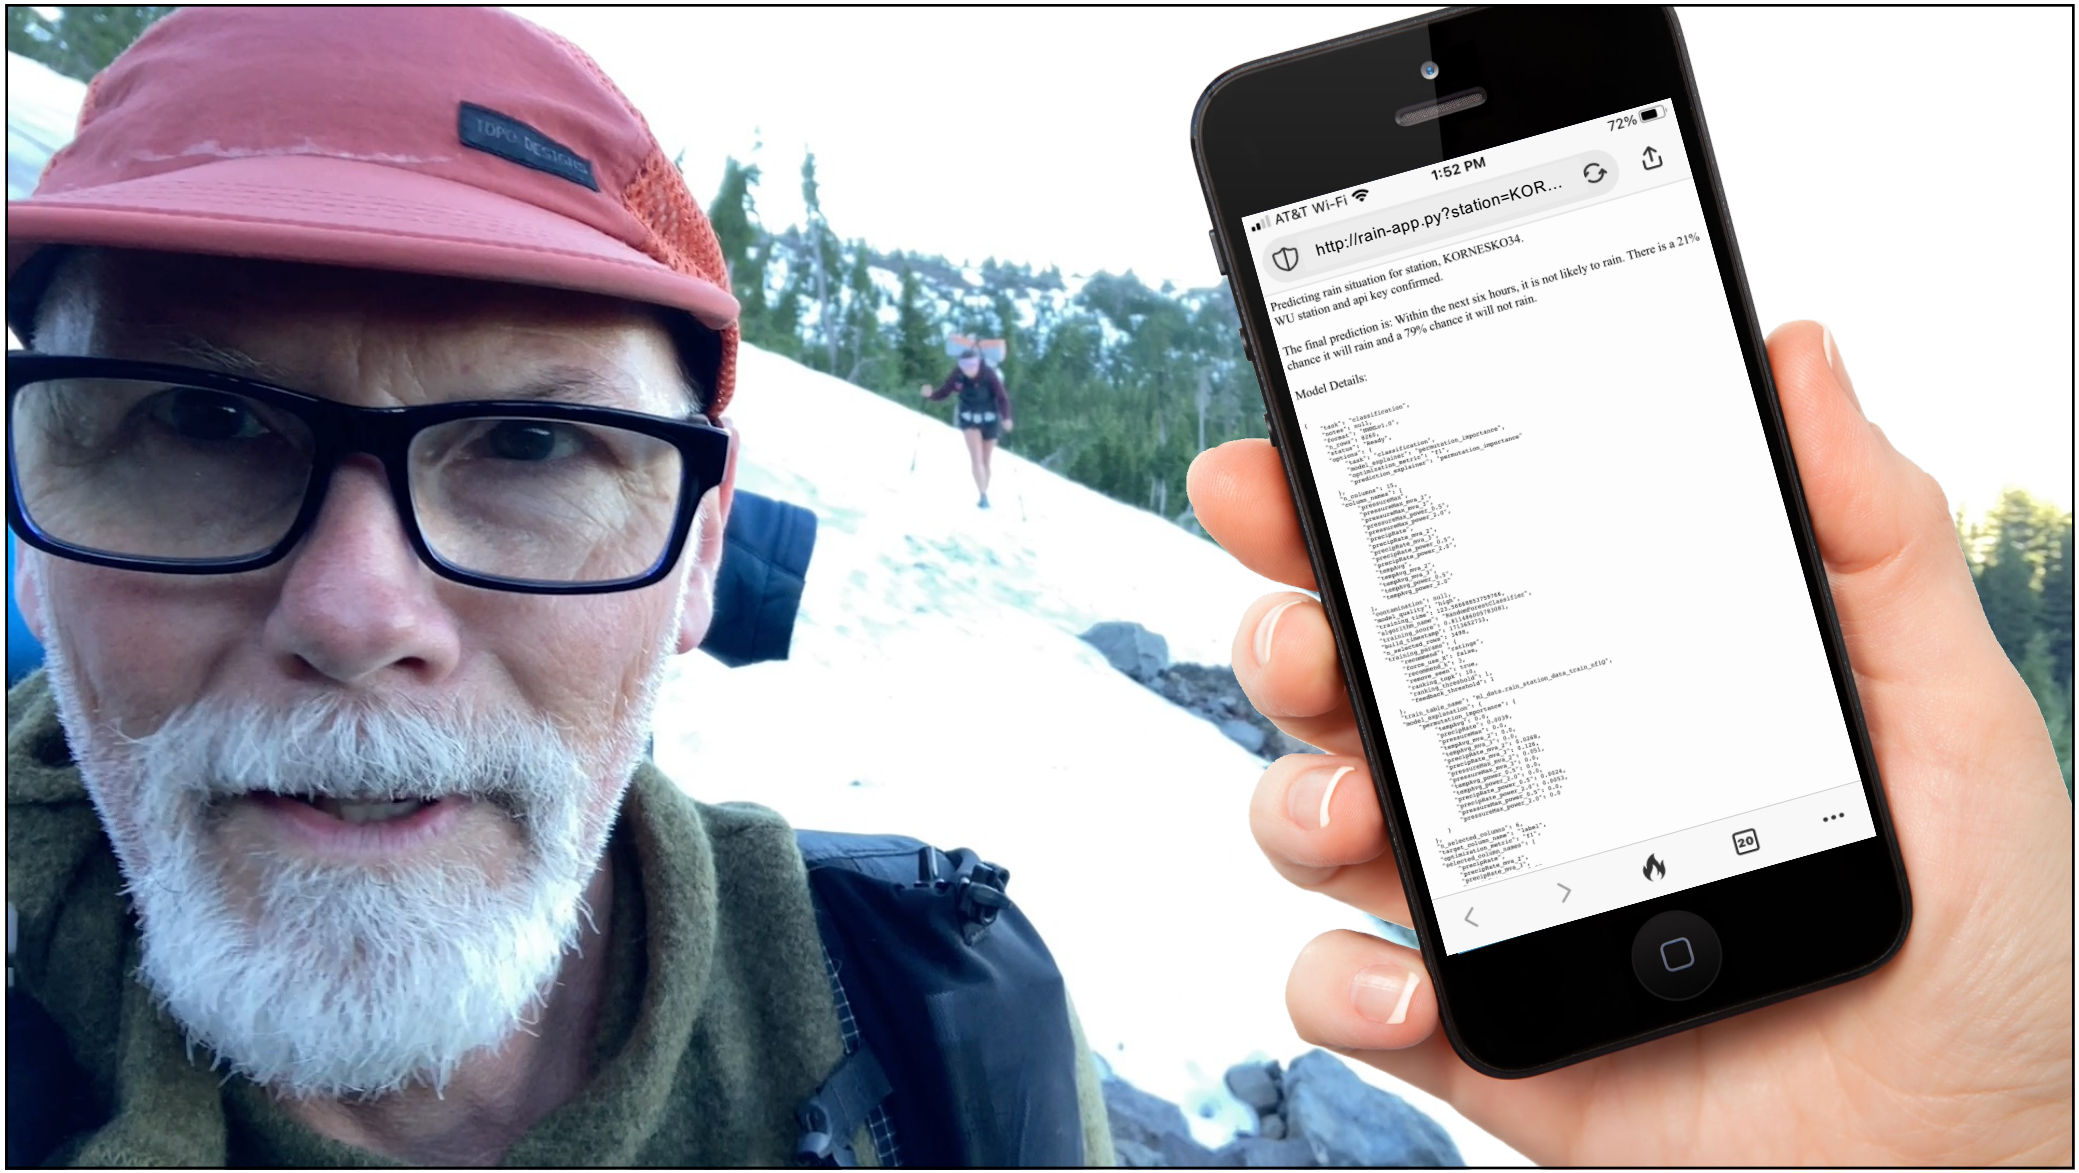
AT&T Wi-Fi
1:52 PM
72%
http://rain-app.py?station=KOR...
Predicting rain situation for station, KORNESKO34.
WU station and api key confirmed.
The final prediction is: Within the next six hours, it is not likely to rain. There is a 21% chance it will rain and a 79% chance it will not rain.
Model Details:

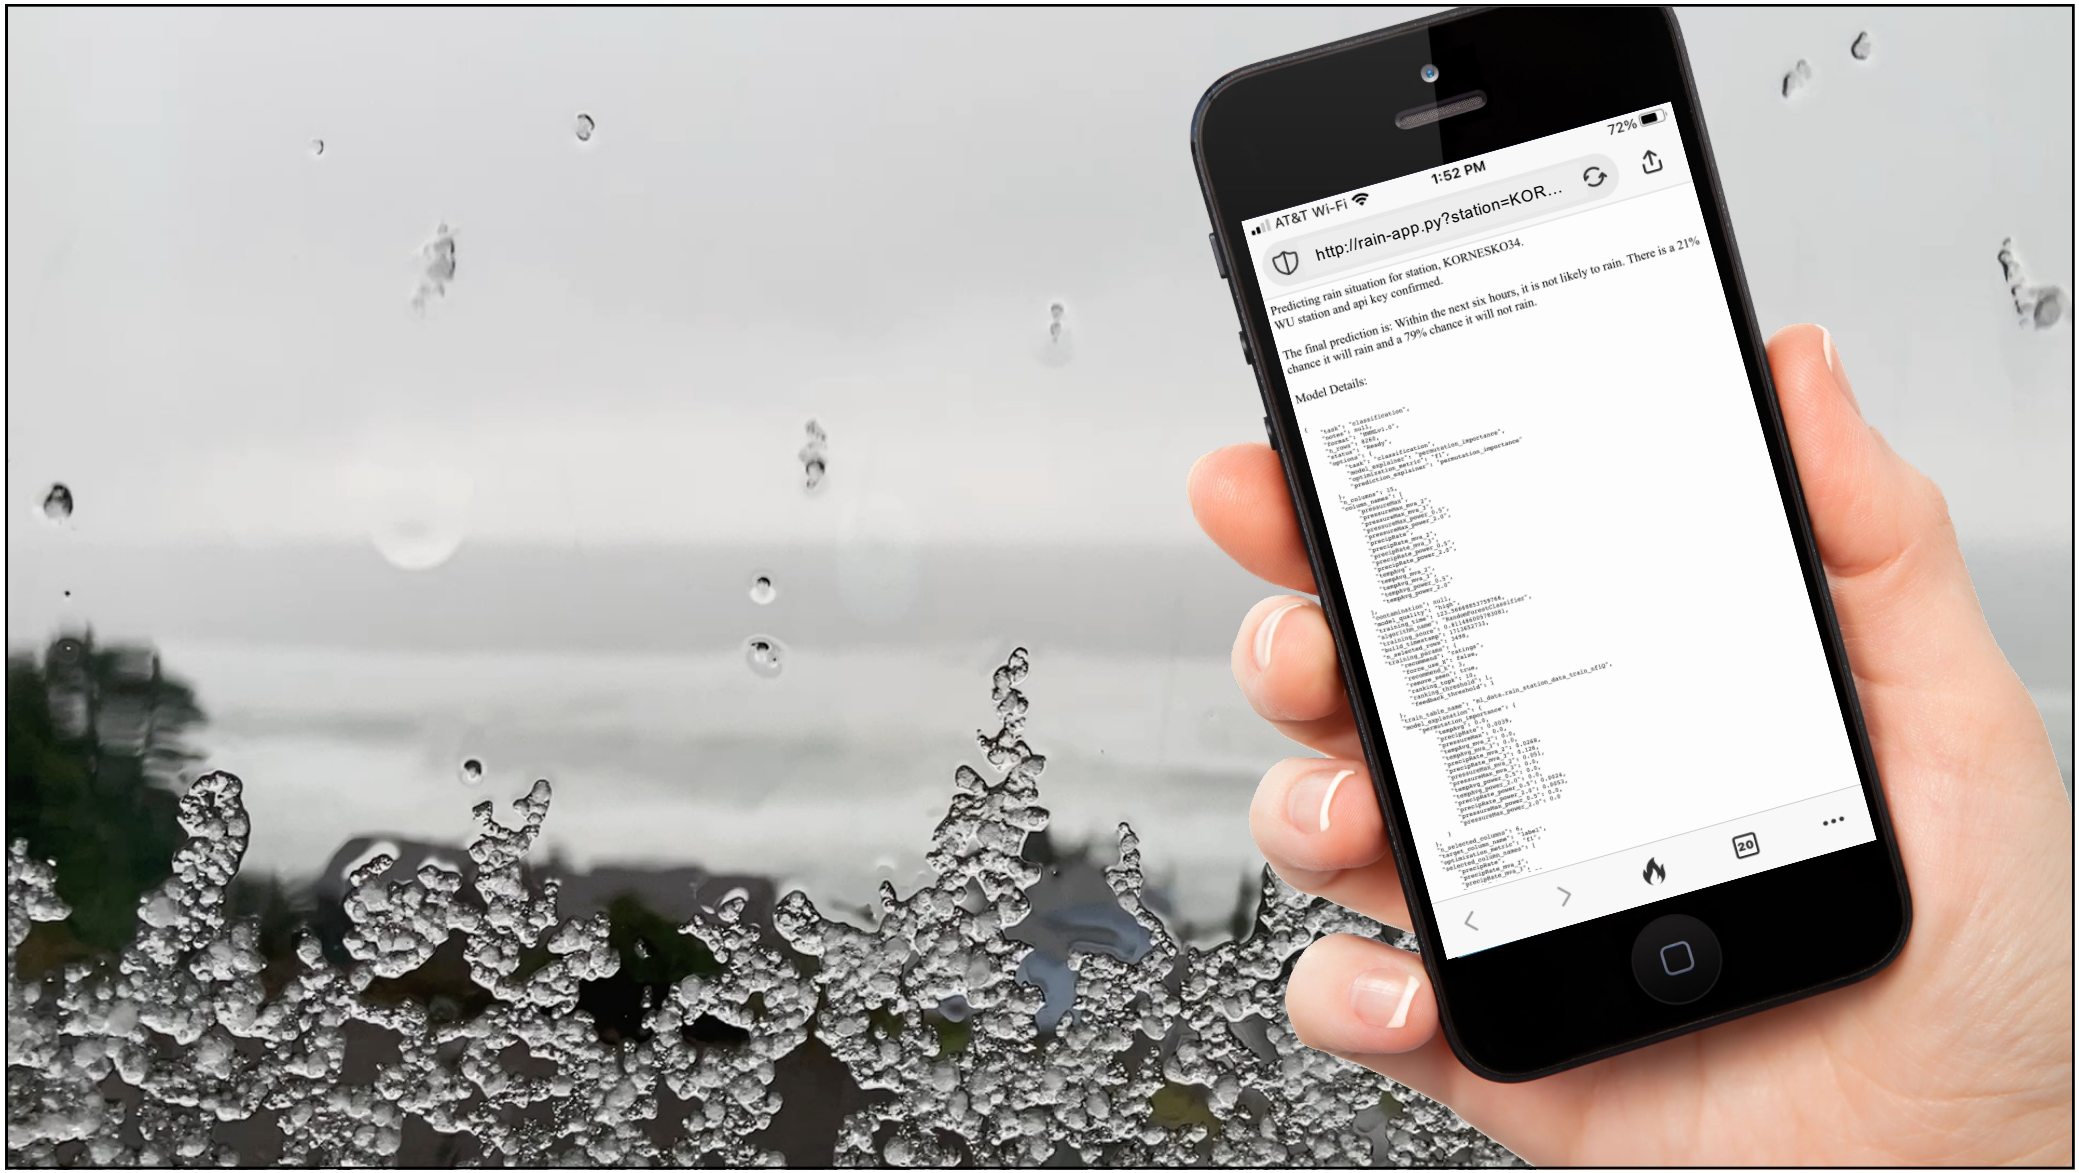
AT&T Wi-Fi
1:52 PM
72%
http://rain-app.py?station=KOR...
Predicting rain situation for station, KORNESKO34.
WU station and api key confirmed.
The final prediction is: Within the next six hours, it is not likely to rain. There is a 21% chance it will rain and a 79% chance it will not rain.
Model Details:

The Situation

It's helpful if I know it's going to rain.

Most of my hikes are less than six hours.

What I need is a six hour rain forecast.

Ever hear of a 6hr rain forecast? No.

And, local forecast aren't really so local.

Perhaps AI can help?

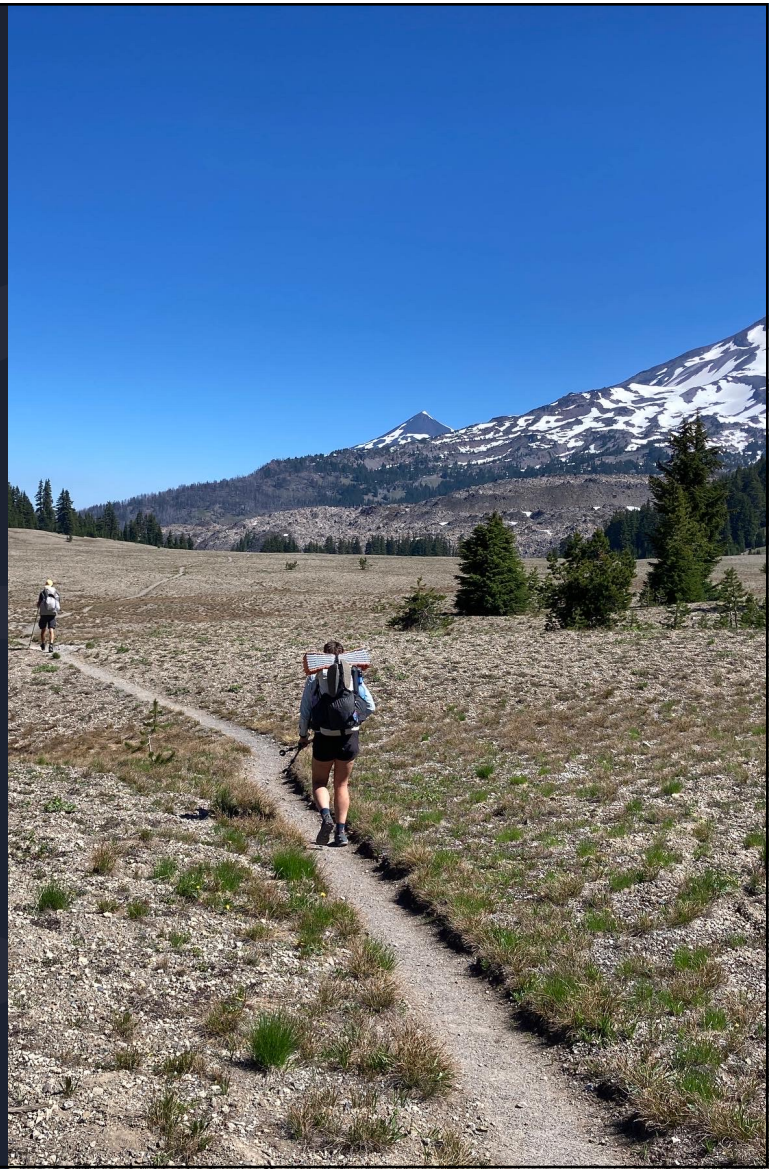

Objective Summ

Build an AI,

to forecast the likelihood of rain,

from now out to six hours into the futu

Needs to run from a URL on my phone.

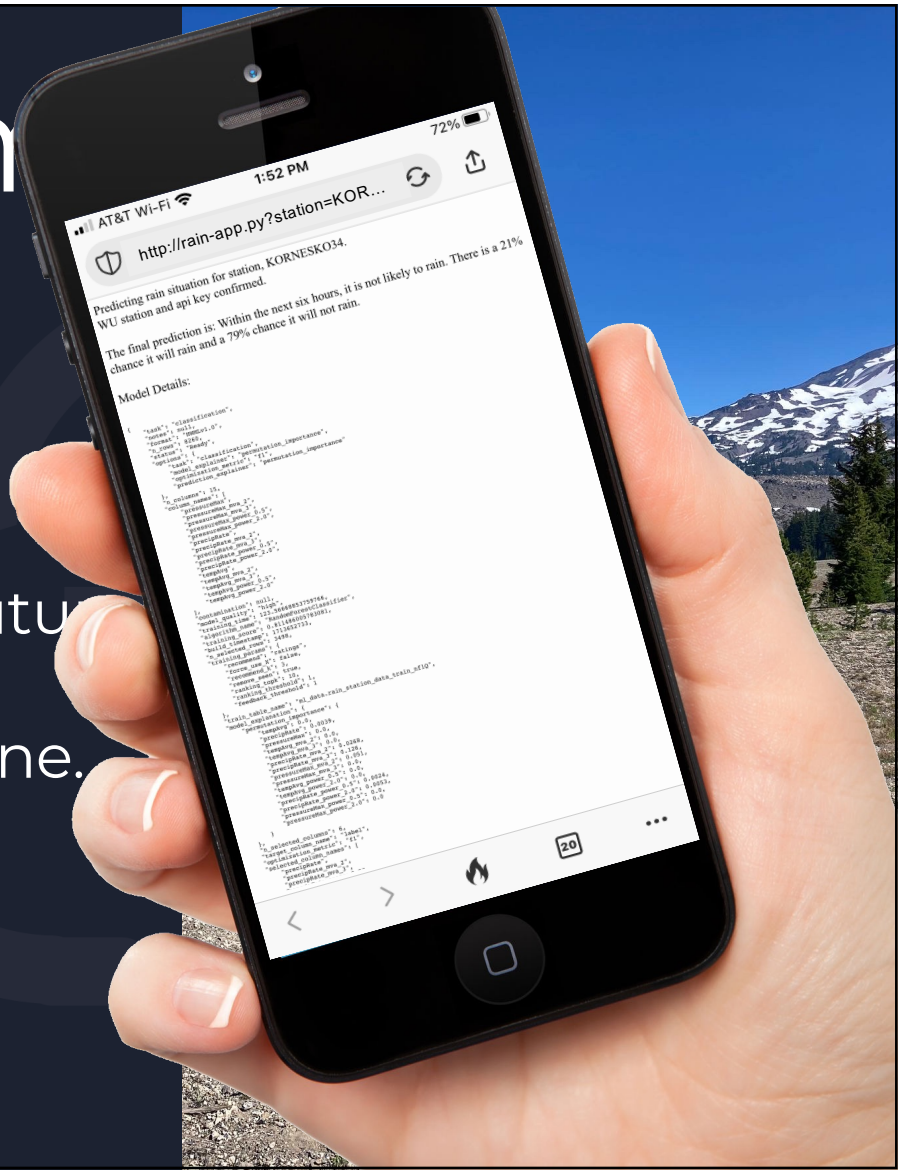

THE FLOW

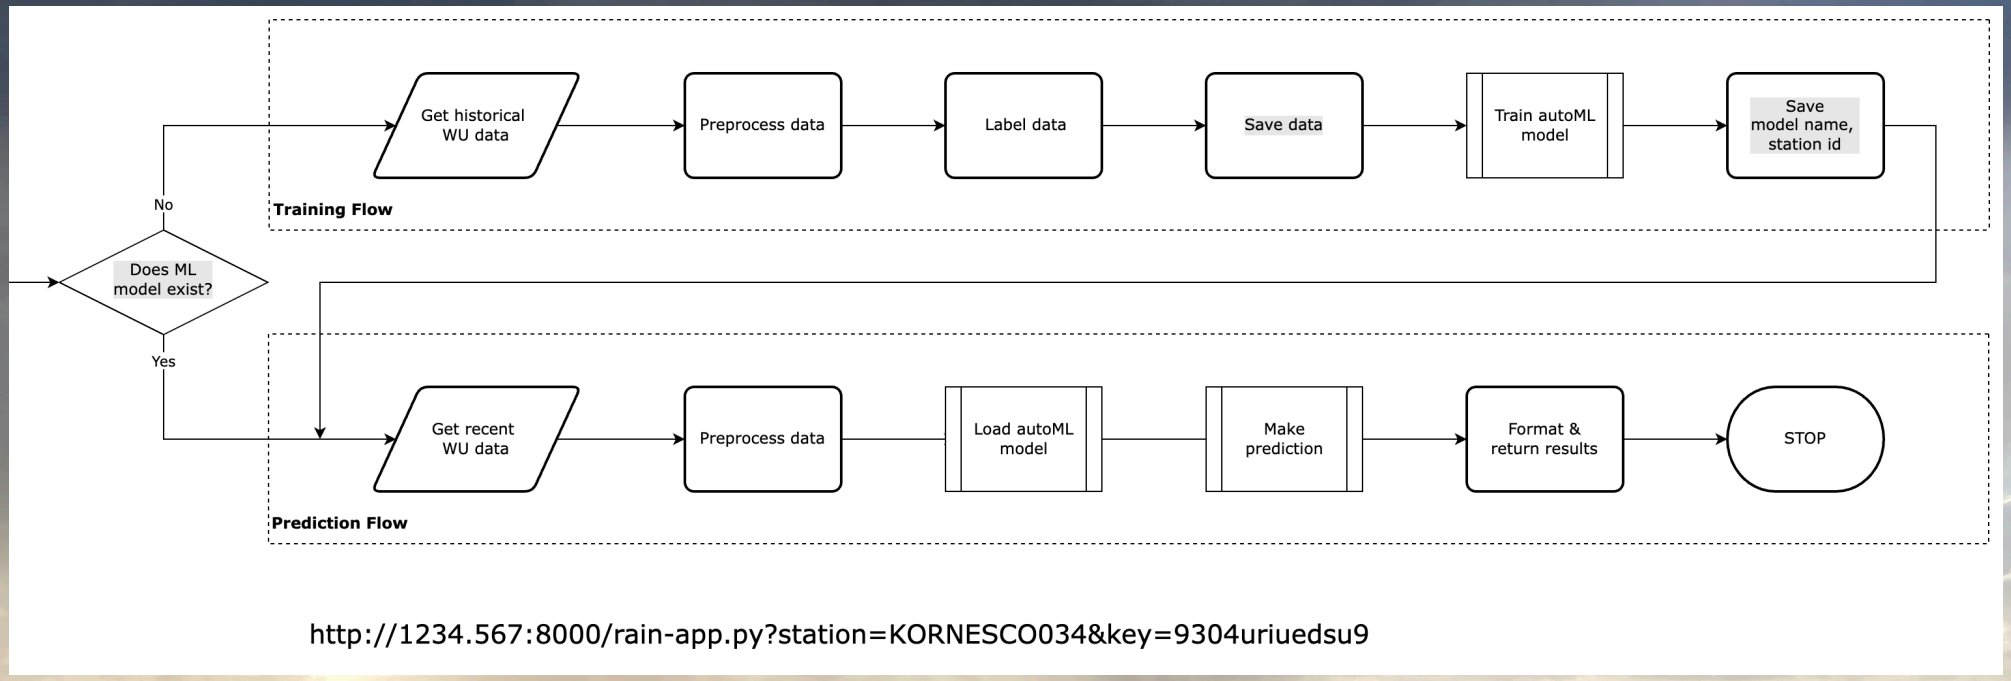

http://1234.567:8000/rain-app.py?station=KORNESCO034&key=9304uriuedsu9

The WU station ID and API key have been validated to ensure weather data can be retrieved.

THE ARCHITECTURE

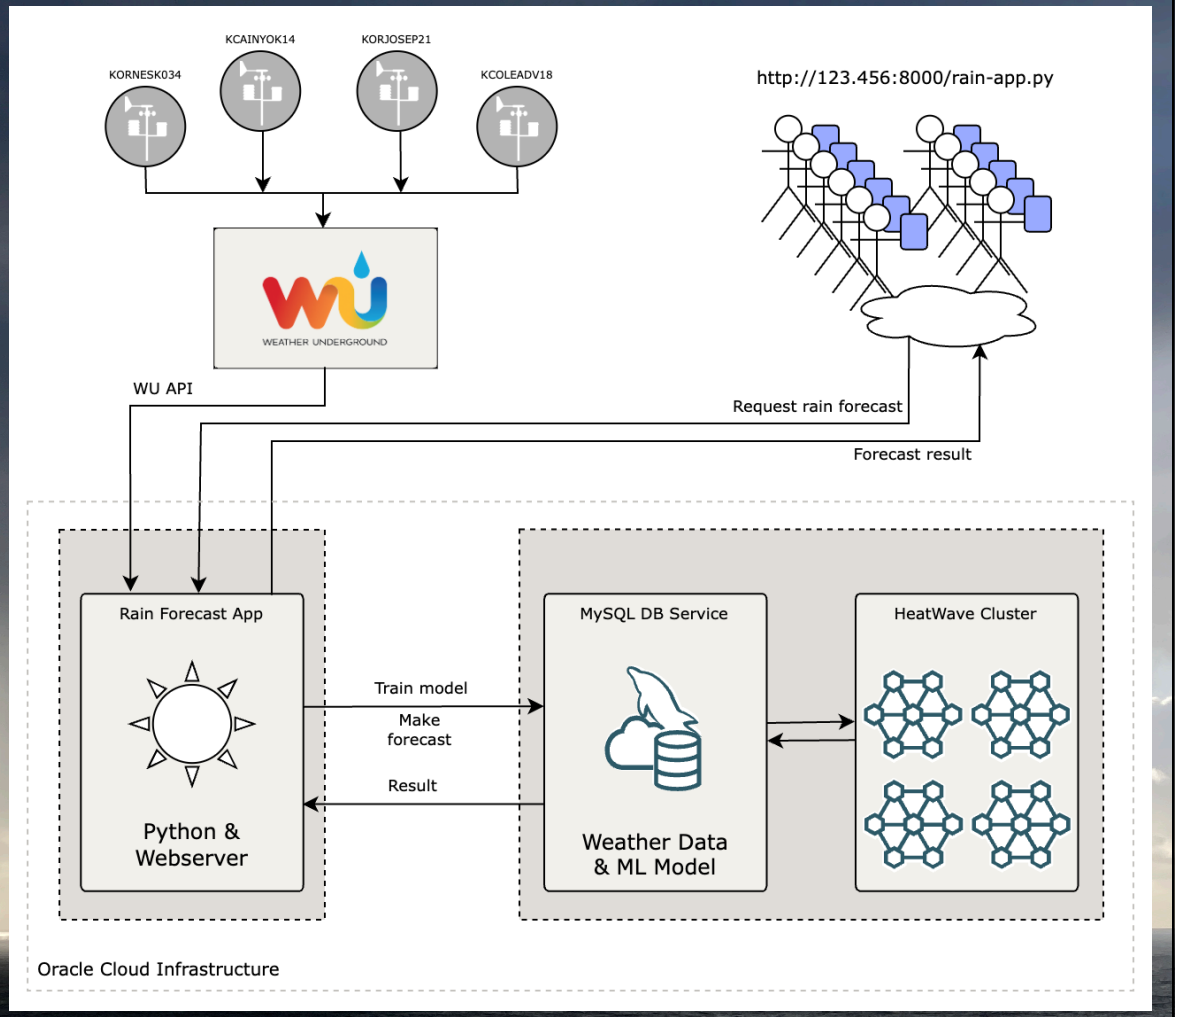

Train using MySQL HeatWave AutoML

```
    # rain_station_data_train_000 must be session specific
    training_data_session = f"ml_data.rain_station_data_train_{session_rnd_str}"
    sql_mysql(sql=f"drop table if exists {training_data_session}", session=session, v=V)
    sql =f"""
    CREATE table {training_data_session}
    AS
    select  label,
            pressureMax, pressureMax_mva_2, pressureMax_mva_3, `pressureMax_power_0.5`, `pressureMax_power_2.0`,
            precipRate, precipRate_mva_2, precipRate_mva_3, `precipRate_power_0.5`, `precipRate_power_2.0`,
            tempAvg, tempAvg_mva_2, tempAvg_mva_3, `tempAvg_power_0.5`, `tempAvg_power_2.0`
    from    rain_station_data where station_id = '{station_id}'
    """
    sql_mysql(sql=sql, session=session, v=V)

    sql=f"CALL sys.ML_TRAIN('{training_data_session}', 'label',JSON_OBJECT('task', 'classification', 'optimization_metric', 'f1'), @rain_model)"
    sql_mysql(sql=sql, session=session, v=V)
    result = query_mysql(sql="select @rain_model", engine=engine, v=V)
    model_handel = str(result.values[0,0])
    logit(f"MAIN: TRAIN COMPLETE: model_handel = {model_handel}")
```

It is really this simple. Create a table with your preprocessed data, set the sys.ML_TRAIN parameters and run sys.ML_TRAIN just like any other MySQL procedure.

Predictions using MySQL HeatWave

```
def make_prediction(station_id, wu_api_key, engine, session, v:bool=False):

    try:
        if v: logit(f"make_prediction. Entering. station_id:{station_id}")

        # Check is model exists
        sql=f"select model_handle from rain_ml_models where station_id = '{station_id}'"
        result = query_mysql(sql=sql, engine=engine, v=v)
        model_handle = str(result.values[0,0])
```

First, I need to check if the ML model has already been created. So, I make a simple MySQL query to my rain_ml_models table.

Predictions using MySQL HeatWave

```
# Get most recent row of preprocessed [illegible] prediction data @row_input
most_recent_data = df.tail(1)
if v: logit(f"make_prediction. Most recent data: df.tail(1)\n{df.tail(1)}")

# Used column_names list to ensure I build @row_input correctly
sql=f"select column_names from ML_SCHEMA_admin.MODEL_CATALOG where model_handle = '{model_handle}'"
result = query_mysql(sql=sql, engine=engine, v=v)
column_names = str(result.values[0,0])

if v: logit(f"make_prediction. column_names:{column_names}")
if v: logit(f'make_prediction. most_recent_data["pressureMax"].values[0]:{most_recent_data["pressureM
```

To make a prediction using an existing ML model, the most recent data is given to the ML model to make the prediction.

I have lots of preprocessed columns (most_recent_data) so I need to ensure I give the model only the columns (column_names) it wants.

Predictions using MySQL HeatWave

```
recent_row = f"""
SET @row_input = JSON_OBJECT(
'pressureMax', {most_recent_data['pressureMax'].values[0]},
'pressureMax_mva_2', {most_recent_data['pressureMax_mva_2'].values
'pressureMax_mva_3', {most_recent_data['pressureMax_mva_3'].values
'pressureMax_power_0.5', {most_recent_data['pressureMax_power_0.5
'pressureMax_power_2.0', {most_recent_data['pressureMax_power_2.0
'precipRate', {most_recent_data['precipRate'].values[0]},
'precipRate_mva_2', {most_recent_data['precipRate_mva_2'].values[
'precipRate_mva_3', {most_recent_data['precipRate_mva_3'].values[
'precipRate_power_0.5', {most_recent_data['precipRate_power_0.5'].values[0]},
'precipRate_power_2.0', {most_recent_data['precipRate_power_2.0'].values[0]},
'temp
'temp
'temp
'tempAvg_power_0.5', {most_recent_data['tempAvg_power_0.5'].values[0]},
'tempAvg_power_2.0', {most_recent_data['tempAvg_power_2.0'].values[0]})
"""
if v: logit(f"make_prediction. Creating SET @row_input:\n{recent_row}\n")
sql_mysql(sql=recent_row, session=session, v=v)
```

The most recent weather data is placed into the @row_input session variable. @row_input will be given to the trained ML model to help make the prediction.

Predictions using MySQL HeatWave

```
# Load ML model into HeatWave
if v: logit(f"make_prediction. Loading model_handle:{model_handle}")
sql_mysql(sql=f"CALL sys.ML_MODEL_LOAD('{model_handle}', NULL)", session=session, v=v)

# Make prediction
sql = f"SELECT sys.ML_PREDICT_ROW(@row_input, '{model_handle}', NULL)"
if v: logit(f"make_prediction. Making Prediction, sql={sql}")
result = query_mysql(sql=sql, engine=engine, v=v)
result = str(result.values[0,0])
if v: logit(f"make_prediction. result={result}\n\ntype(result):{type(result)}")
```

Before I can request a prediction, the trained ML model must be "loaded" into HeatWave cluster.

Once loaded, the model_handle and the @row_input is given to make the prediction... using a simple SELECT statement!

Predictions using MySQL HeatWave

```
# Unload ML model from HeatWave
if v: logit(f"make_prediction. UN loading model_handle:{model_handle}")
sql_mysql(sql=f"CALL sys.ML_MODEL_UNLOAD('{model_handle}')", session=session, v=v)
```

After the quick prediction occurs, the trained ML model is "unloaded" from the HeatWave cluster.

Then the results are sent to the user's mobile device!

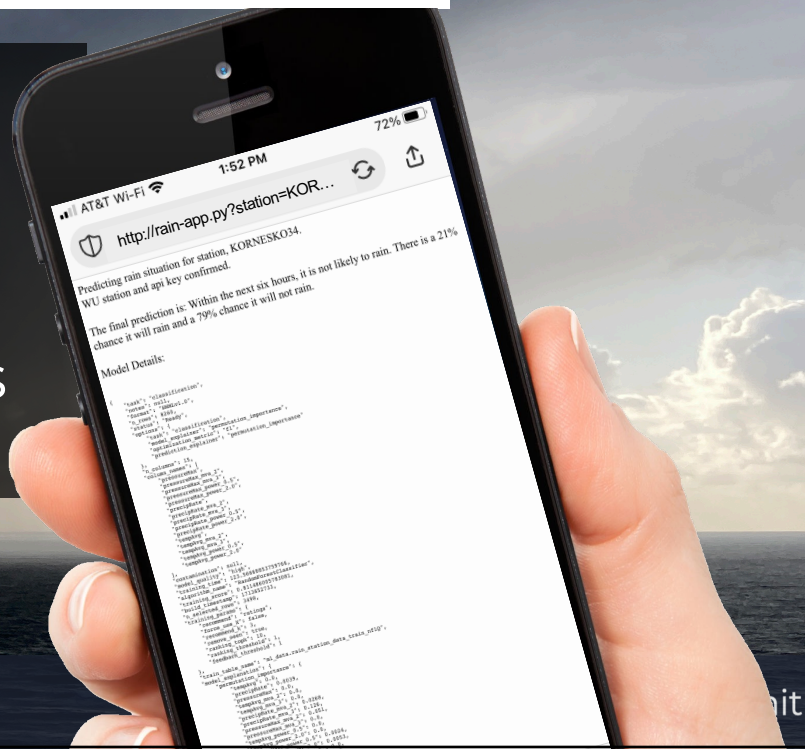

SURPRISES – IT JUST WORKS

- **AutoML** quality, simplicity, HeatWave speed, usability and especially stability.
- **Documentation** is created to help me use the product.
- **Integration (MySQL, HeatWave)** with Python means APIs anywhere, community everywhere and very straightforward.

A special shout-out to both

Kirk Bohn & Preside Foster

Thank you both for helping make this presentation possible!


Now it's your turn!
www.orapub.com for presentation and code.
AT&T Wi-Fi
1:52 PM
72%
http://rain-app.py?station=KOR...
Predicting rain situation for station, KORNESKO34.
WU station and api key confirmed.
The final prediction is: Within the next six hours, it is not likely to rain. There is a 21% chance it will rain and a 79% chance it will not rain.
Model Details:
viscosityna.com
@ViscosityNA

Thank You!

ORAPUB
POWERED BY VISCOSITY

VISCOSITY

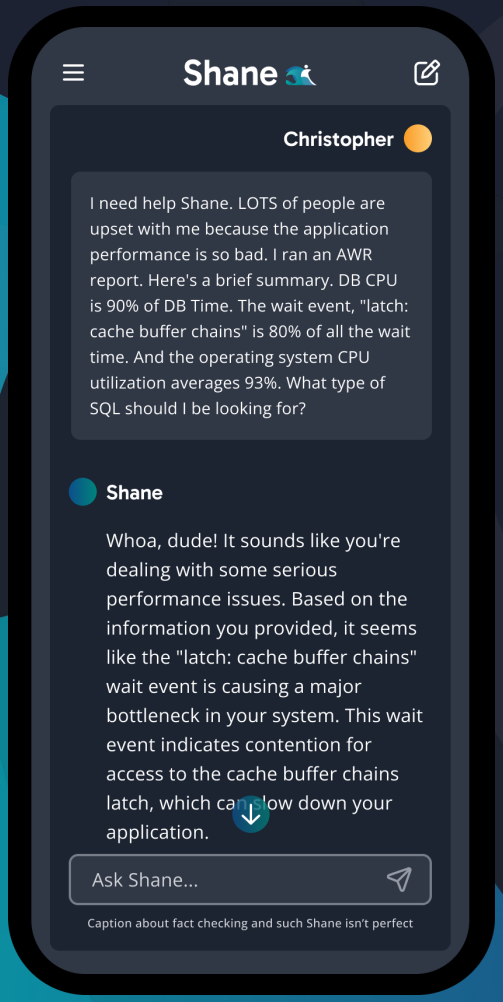

viscosityna.c

ve Summit

VISCOSITY
PROFESSIONAL SERVICES

Crafting the Ultimate Outdoor Adventure Forecast Using MySQL HeatWave

MySQL and HeatWave Summit 2024 | May 1, 2024
Oracle Conference Center, Redwood City, California USA

MySQL HeatWave

ORACLE

Shane AI

Craig Shallahamer

Oracle ACE Director
Applied AI Scientist | OraPub Founder

@orapub
linkedin.com/in/craig-shallahamer-571a94a/
craig.shallahamer@viscosityna.com