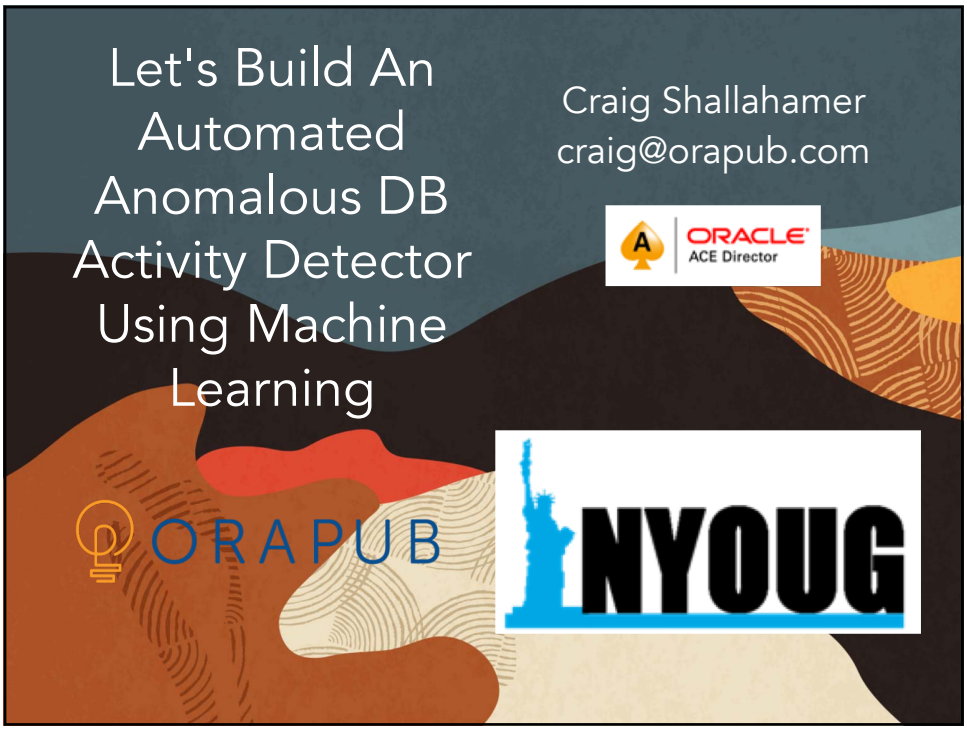

This presentation was given by Craig Shallahamer (craig@orapub.com) at the October 13, 2020 NYOUG Conference. There is likely a more recent version at www.orapub.com > Login > Presentations

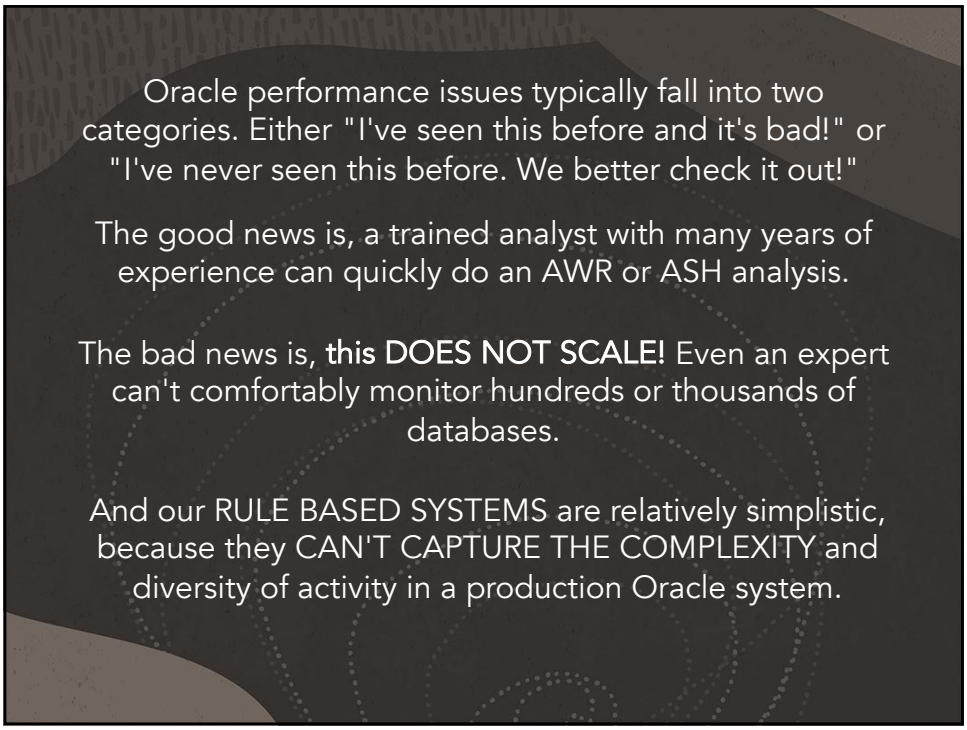

One solution for this unsustainable monitoring and analysis problem is to use machine learning.

How to apply machine learning to quickly and automatically detect an anomalous or recognized performance situations (all before the phone rings, is what this presentation is all about.

One solution for this unsustainable monitoring and analysis problem is to use machine learning.

Besides an introduction to machine learning in the Oracle DBA world, how to apply machine learning to quickly and automatically detect an anomalous performance situation, all before the phone rings, is what this presentation is all about.

Anomalous performance
does not imply poor performance.

About Me...

- Long time Oracle DBA
- Specialize in **predictive analytics, machine learning** and Oracle performance tuning
- Performance researcher
- Blogger: A Wider View About Oracle Performance Tuning
- Author: Oracle Performance Firefighting and Forecasting Oracle Performance.
- Conference speaker
- Teacher and mentor
- Oracle ACE Director
- Quest/IOUG DBA Track Manager

ORAPUB

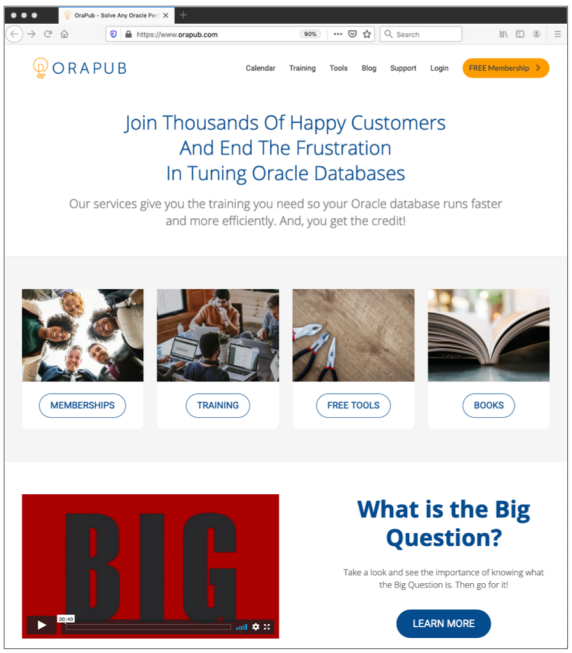

OraPub works with IT to deploy machine learning into their monitoring and alerting processes.

OraPub works with Oracle DBAs empowering them to beat bots, AI, machine learning and autonomous anything.

ORAPUB

Want my updated slides?

`www.orapub.com >`
`Login >`
`Tools & Presentations`

ORAPUB

My Machine Learning Journey

It all started in 1990

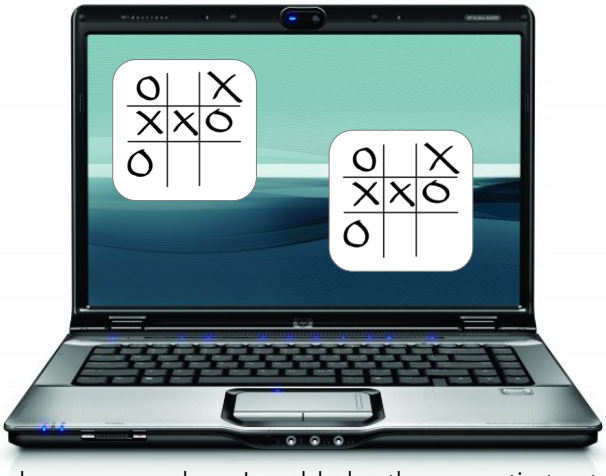

I created a program where I could play the game tic-tac-toe with a computer, but also enabled the computer to play itself… and learning through that process.

ORAPUB

In 1997 Neural Networks…

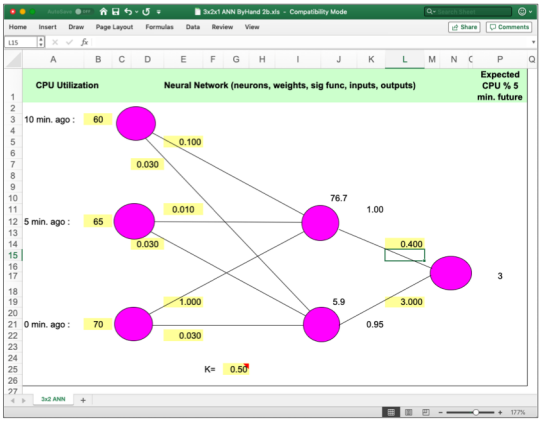

I was obsessed is performance prediction… forecasting what was likely to happen or not happen in the future.

ORAPUB

In 2001 at RMOUG

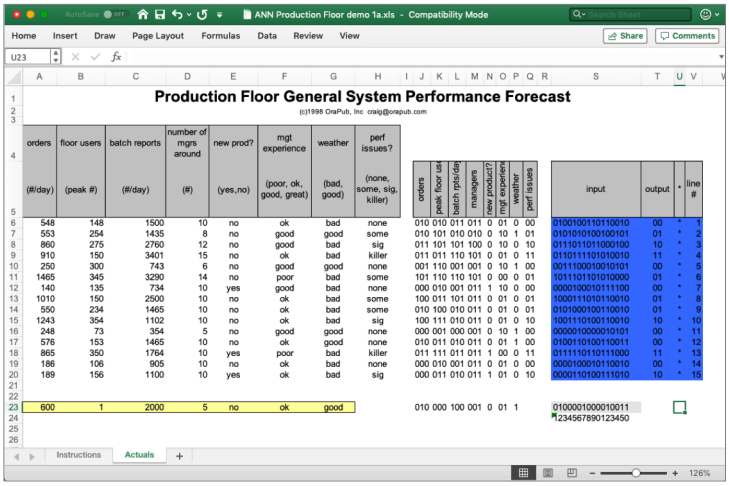

In 2001, after two years of rejections, I presented on ANNs at the RMOUG conference. My goal was to predict poor performance in the near future. While ANN is sound, I was never able to realize my "near future" dream. But I believe that dream could now be realized.

ORAPUB

2007

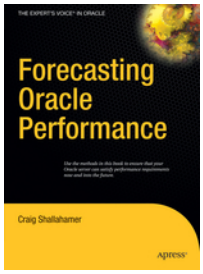

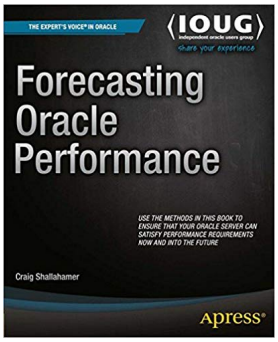

ORAPUB

What Is Machine Learning

What Is Machine Learning?

- ML fits under the umbrella of AI.
- At it's core, ML is about understanding data; extracting interesting and useful patterns. But this is done **methodically** and using a wide variety of **algorithms**.
- ML contains a growing **set of algorithms** to analyze data. Here is a short list: Support Vector Machines (SVM), Decision Tree Learning, Instance-Based Learning, Generalized Linear Models, Artificial Neural Network, Centroid-Based Clustering, Hierarchical Clustering, Density-Based Clustering.
- ML involves using a variety of advanced statistical and computing **techniques** to process data to find patterns; feature selection, feature engineering, imputation, stratification, principle component analysis, cross fold validation, residual analysis, data transformation, centering and scaling, etc.

ORAPUB

Why All The Buzz Now?

- **Management looking to implement AI** to demonstrate they are cutting edge, reduce repetitive and timing consuming tasks, and do things their competition can't.
- **Software**. Powerful, "free" and available to anyone. It's available on Windows, Mac OS X and Linux. And, the serious number crunching can be run in parallel!
- **Python** and **R** are free statistical packages. Both have powerful libraries, enabling a wide variety of machine learning tasks. Did I forget to mention they are free?
- **Oracle & ML**. Move the algorithms; Not the Data!
  - A tremendous amount of time and effort is spent simply moving data around so it can be processed.
  - All the Oracle Database data and transaction management is, of course, available.
  - Oracle is moving more and more ML processes directly into Oracle, including the Autonomous Database. These are and will be accessible through SQL commands and PL/SQL procedures.
- **Computing Power**. CPUs and GPUs (fast matrix multiplications) are crazy fast and ML software is designed to take advantage of this power. Processing can be in a cloud.
- **Lots of data** is now available. Repeat.
- **Organizational Impact**. ML is touching all business aspects in new ways… IT !
- **Domain Expertise**. **Oracle performance analysts have a unique advantage and place in this movement.**

ORAPUB

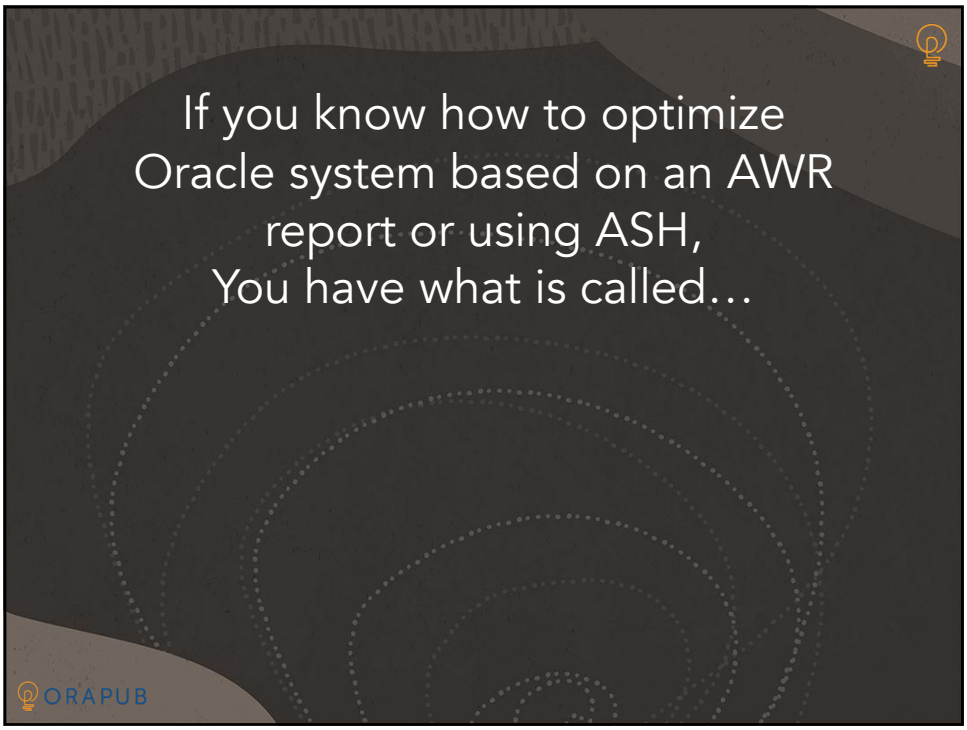
If you know how to optimize
Oracle system based on an AWR
report or using ASH,
You have what is called…
ORAPUB

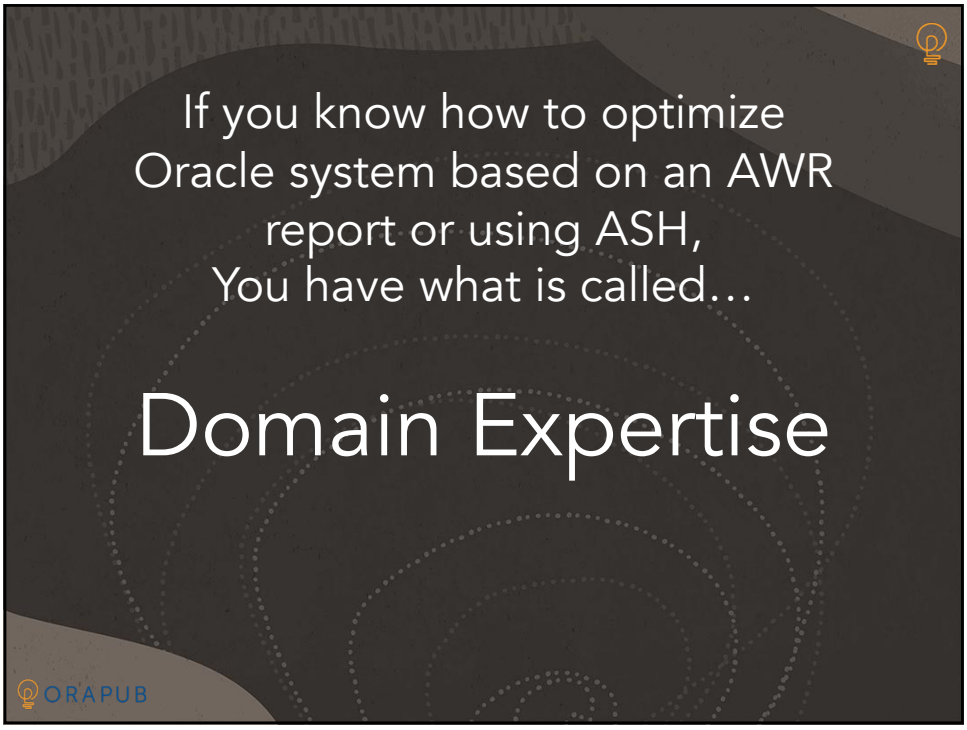
If you know how to optimize
Oracle system based on an AWR
report or using ASH,
You have what is called…
Domain Expertise
ORAPUB

Let's do some Machine learning

"I've seen this before!"

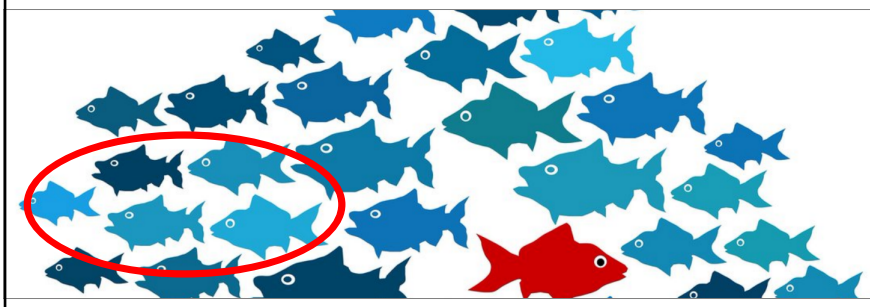

We call this, "classification."

ORAPUB

"I've never seen this before!"

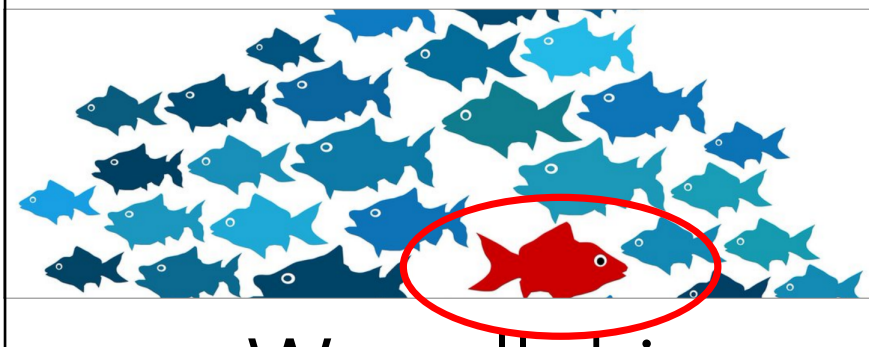

We call this,
"anomaly detection."

ORAPUB

Supervised
Vs
Un-Supervised

Supervised and Unsupervised Learning

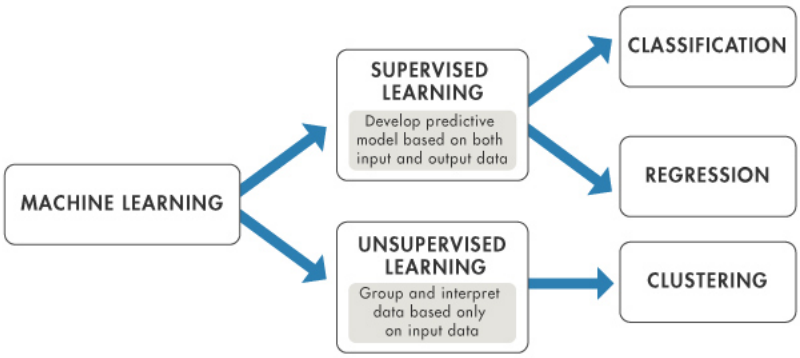

ORAPUB

This presentation was given by Craig Shallahamer (craig@orapub.com) at the October 13, 2020 NYOUG Conference. There is likely a more recent version at www.orapub.com > Login > Presentations

Supervised Vs Un-Supervised

Label

snap_id	uc_psec	aas	trx_psec	perf
1001	2500	34.25	9.45	bad
1002	1200	14.50	6.50	good
1003	1150	16.50	16.50	good
1004	1250	18.50	9.50	good
1005	1300	24.50	2.50	bad

Supervised

Will snap_id 1005 result in "the phone ringing?"

snap_id	uc_psec	aas	trx_psec
1001	2500	34.25	9.45
1002	1200	14.50	6.50
1003	1150	16.50	16.50
1004	1250	18.50	9.50
1005	1300	24.50	2.50

Un-Supervised

Is snap_id 1005 anomalous?

ORAPUB

Un-Supervised Learning With One Cluster

"I've *never* seen this before!"

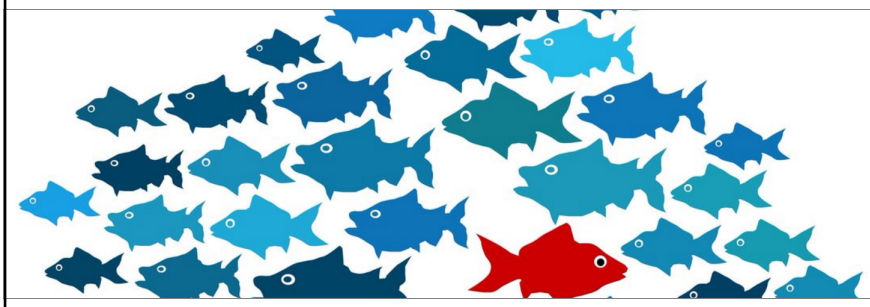

We call this, "anomaly detection."

ORAPUB

Anomaly Detection

(I've never seen this before!)

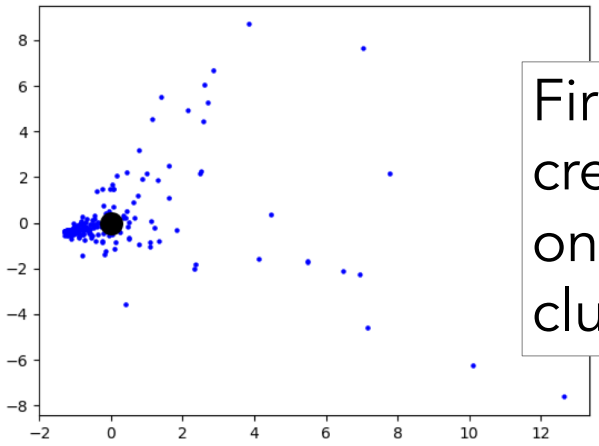

First, create a one class cluster.

ORAPUB

Anomaly Detection

(I've never seen this before!)

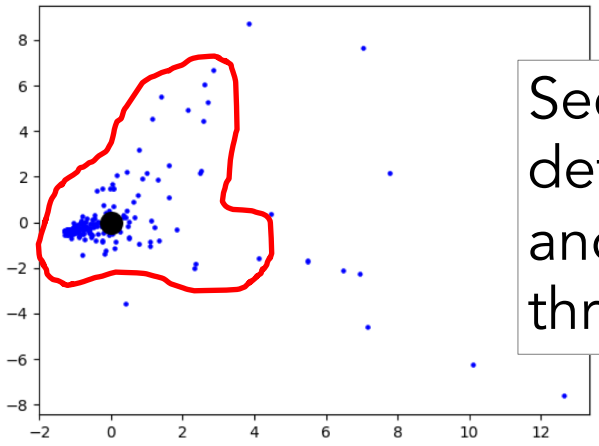

Second, determine anomalous threshold

ORAPUB

Anomaly Detection

(I've never seen this before!)

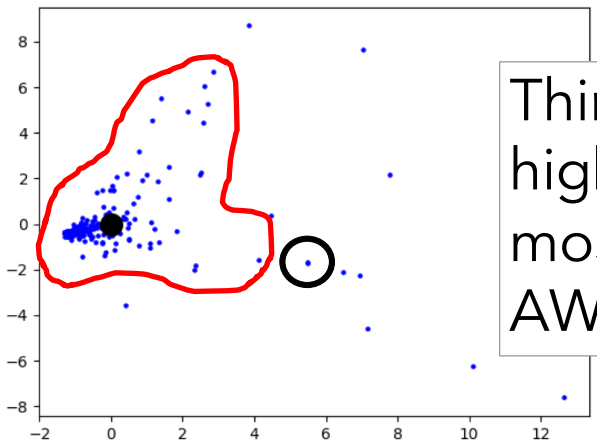

Third, highlight the most recent AWR snap.

ORAPUB

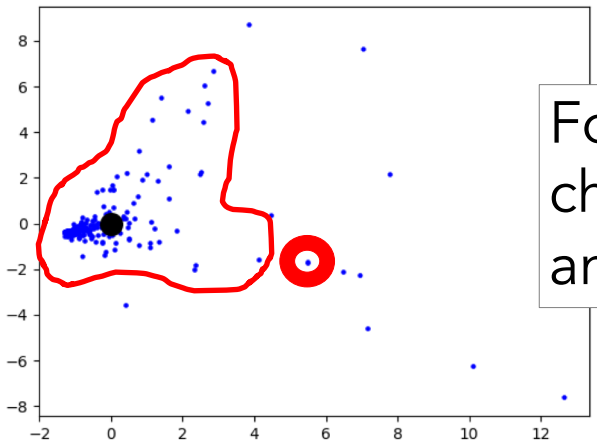

Fourth, check if anomalous

ORAPUB


Let The Typing Begin
To Set Up Your Machine Learning Environment:
blog.orapub.com/20200107/
how-to-setup-your-machine-learning-sandbox-environment.html
You can watch the video on the presentation page:
www.orapub.com > Login > Tools & Presentations


Going Deeper

Resource listing

- **OraPub Membership** **for premium & exclusive content**
  - "How To" Webinars – two each month. Over 110 recorded!!
  - Video Seminars – any device, any time, high quality
  - Learning paths, mentoring, assessments and certificates, priority response
  - 20% LVC discounts
  - You're at the "top of my inbox"
  - SLACK forum exclusively for paid members
- **Live Virtual Classroom (LVC) Training**
  - **Machine Learning For Oracle Professionals**
  - Tuning Oracle Using An AWR Report
  - Tuning Oracle Using Active Session History (ASH) Strategies
  - Core Truths For Oracle Professionals
- **Toolkits** – Many tools available at orapub.com
- **Craig's Blog & Website**
- **Presentations** – www.orapub.com
- **Books:** Oracle Performance Firefighting. Forecasting Oracle Performance.

ORAPUB

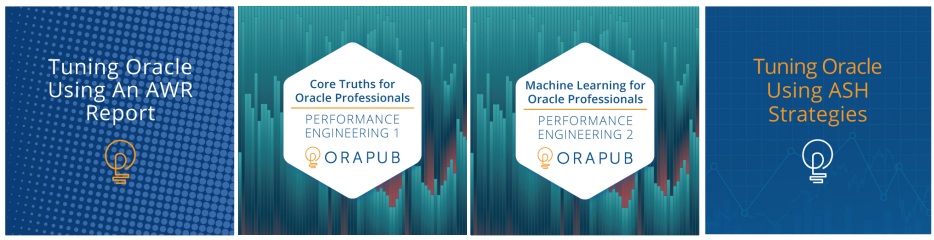

OraPub LVCs Are Different

You learn to master the topic at a deeper, higher confidence and more practical level compared to any other teaching method I offer.

This occurs because the class is spread out over multiple weeks, each session is 2 hours with a day in between, you do homework/activation on your real systems and I personally work with you through the entire class.

Details & Registration: www.orapub.com/lvc
Event Calendar: https://www.orapub.com/events

ORAPUB

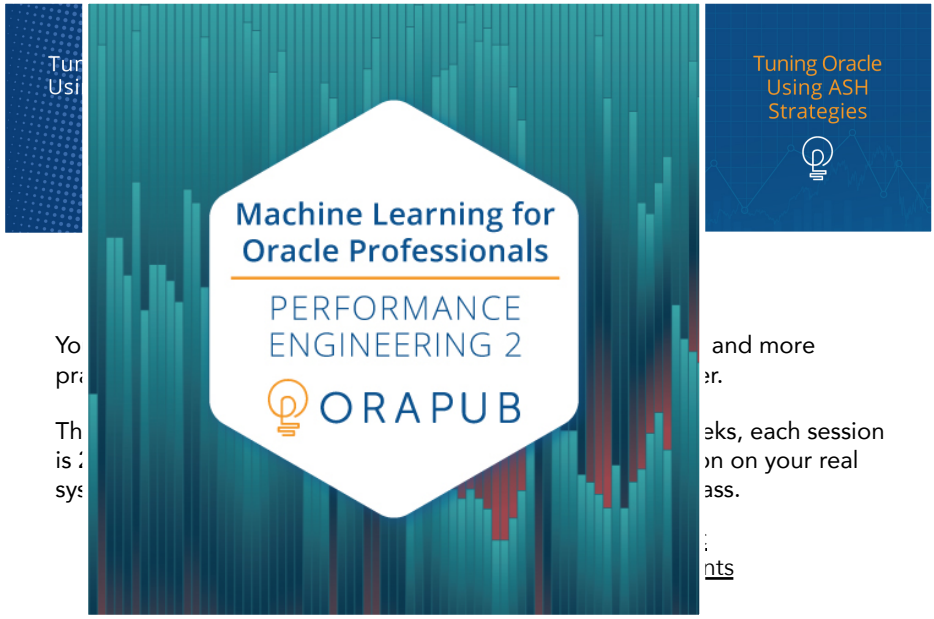

ORAPUB

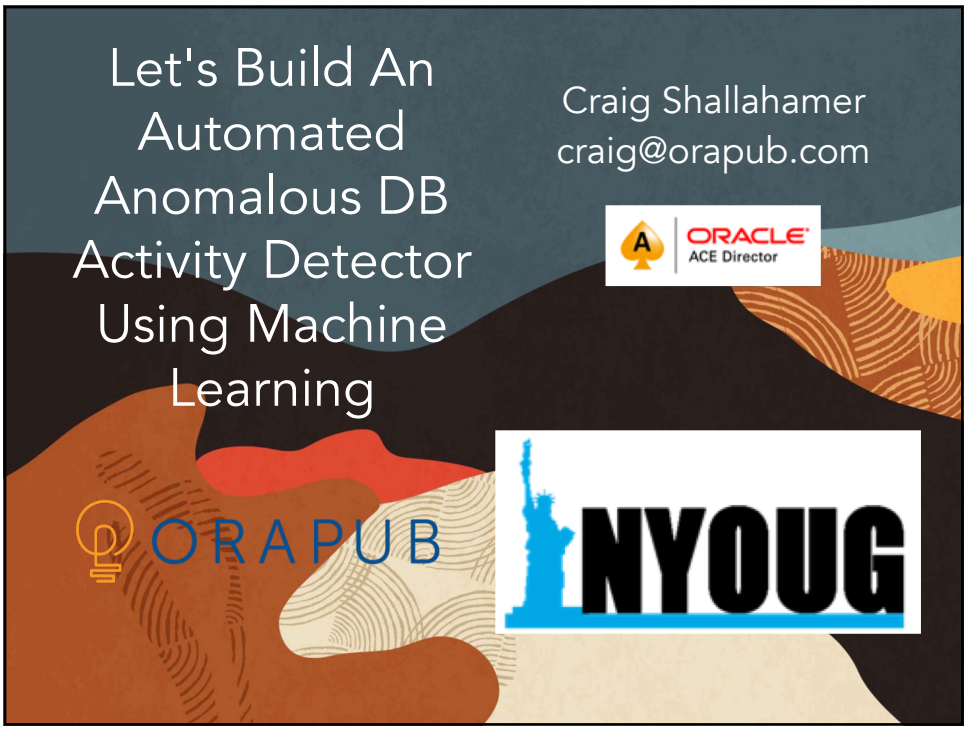
Let's Build An Automated Anomalous DB Activity Detector Using Machine Learning
Craig Shallahamer
craig@orapub.com
ORACLE
ACE Director
ORAPUB
NYOUG