AAS Predictions Using Deep Learning Neural Networks With Sysmetric Data

Craig Shallahamer
craig@orapub.com
craig@viscosityna.com

ORACLE ACE Director

VISCOSITY NORTH AMERICA

ORAPUB
POWERED BY VISCOSITY

NYOUG

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

Realtime Oracle Active Average Sessions Predictions Using Machine Learning

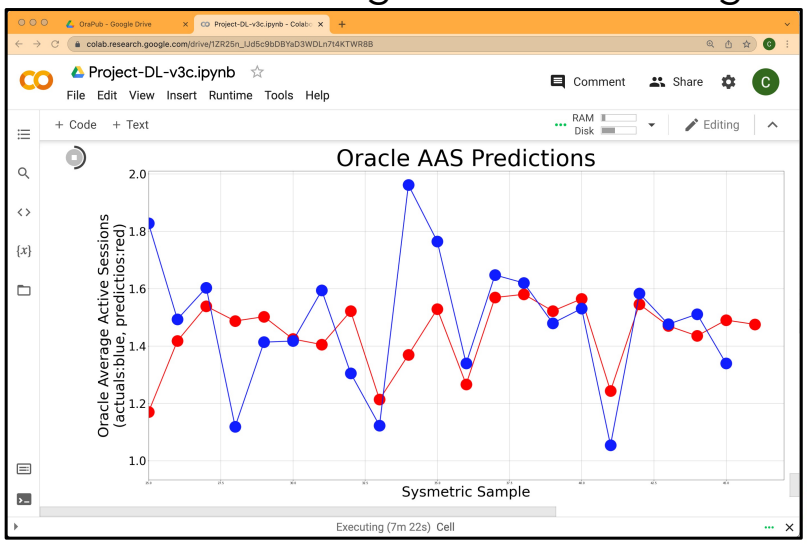

VISCOSITY

NYOUG

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

About Me...

- OraPub's Founder
- Viscosity's Applied AI Scientist
- 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, UTOUG Board Member

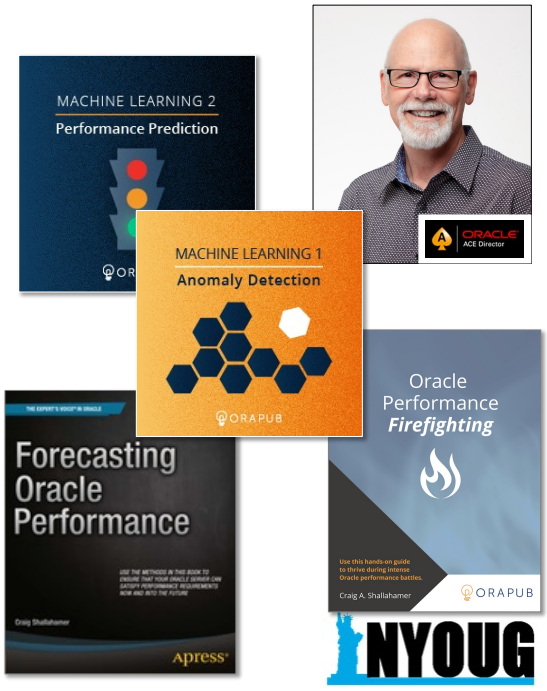

VISCOSITY


VISCOSITY
NORTH AMERICA
Oracle License Management
Get the most out of your Oracle investment
ZERO DOWNTIME
Migrations
Apps
SaaS/PaaS,
Mobility, Application
Development
Professional
Services
Where you need it most
Performance
Health Checks
How's it running?
Staff Aug
Workforce Capacity
on Demand
DBA
Services
Remote and
On-site
On-Call Support
Managed
Services
VISCOSITY
NYOUG


My
Machine Learning
Journey
VISCOSITY
NYOUG

It all started in 1987

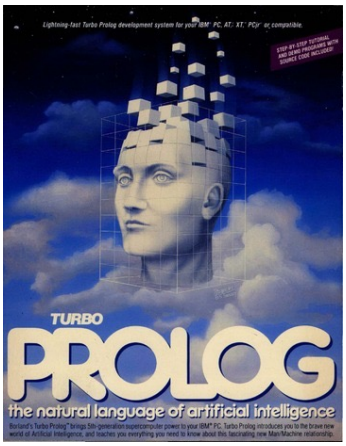

It all started in 1987, when I was intrigued with developing a program that could not only retain data but could answer my questions about that data.

At that time, I was using Borland's *Turbo Pascal* and noticed they released a Prolog product. It was amazing…

I was hooked!

VISCOSITY

NYOUG

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

Then again in 1990

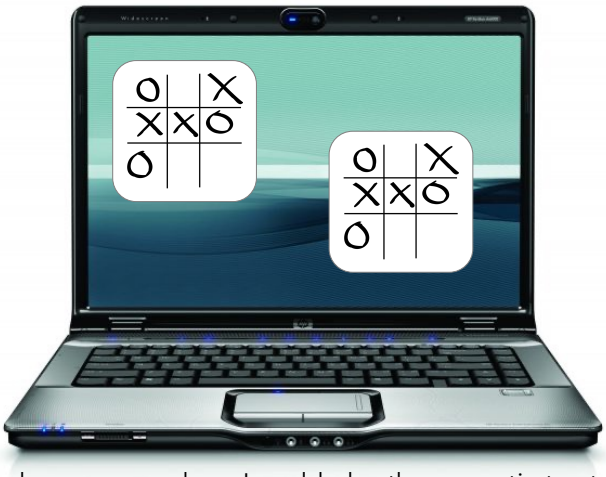

I created a program where I could play the game tic-tac-toe with a computer, but also enabled the computer to play itself… and learning through that process.

VISCOSITY

NYOUG

In 1997 Neural Networks…

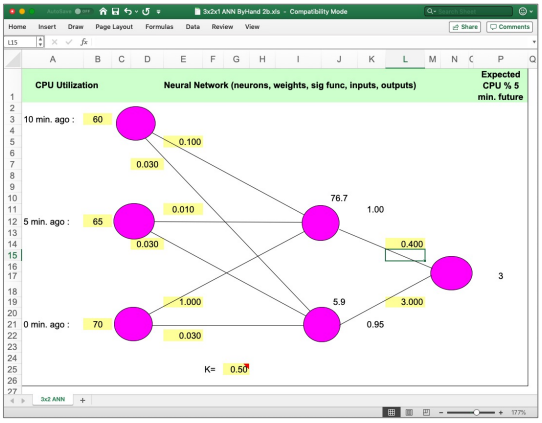

I was obsessed is performance prediction… forecasting what was likely to happen or not happen in the future.

VISCOSITY

NYOUG

In 2001 at RMOUG

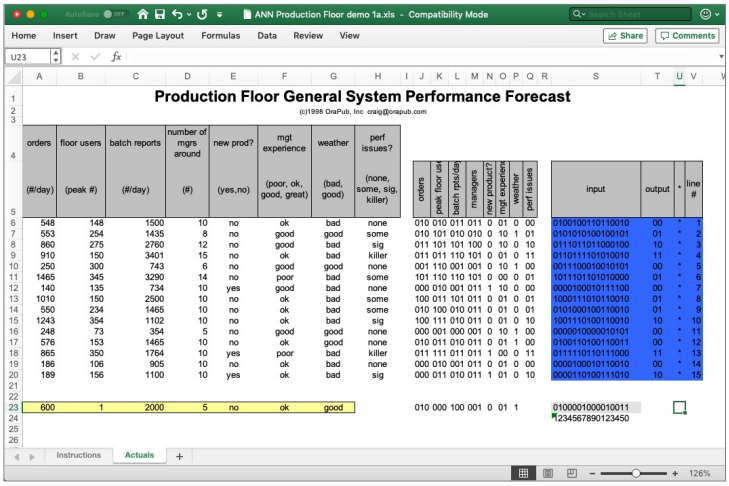

In 2001, after two years of rejections, I presented on ANNs at the RMOUG conference. My goal was to predict poor performance in the near future. While ANN is sound, I was never able to realize my "near future" dream. But I believe that dream could now be realized.

VISCOSITY

NYOUG

2007

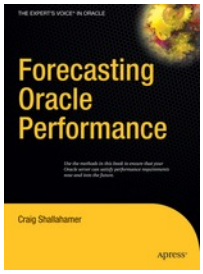

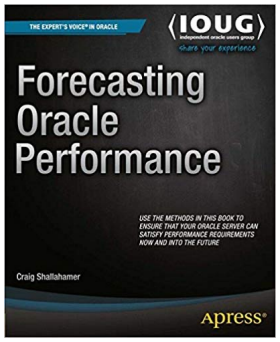

VISCOSITY

NYOUG

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

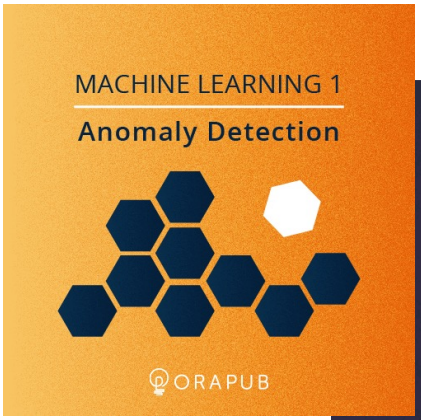

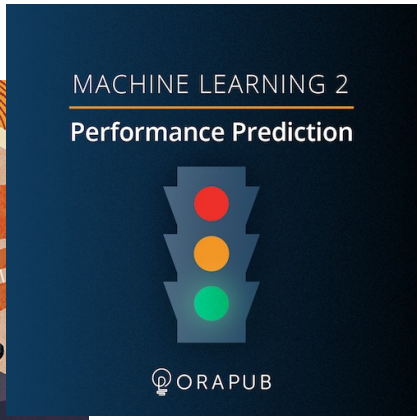

VISCOSITY

NYOUG

Session Focus Areas

- ML use cases for IT
- Deep Learning introduction
- **Project** - Build a real-time graphical numeric AAS prediction system using Deep Learning algorithms, with multivariate dba_hist_sysmetric_summary data.
  - Applied ML process
  - Time-Series Data, a significant challenge
  - ML Essentials
    - Scaling
    - Stepping
    - Feature Selection

VISCOSITY

NYOUG

Uses Cases & Problem Types

VISCOSITY

NYOUG

What Are The Use Cases?

"ML eats rules and code for lunch."

- **Better monitoring and alerting.**
- **Draw our attention to what is important.**
- Self learning and adjusting "rules."
- Better rules… that we can't comprehend or even imagine.
- Manually adjusting and adding new rules does not scale, without ML
- More systems to monitor doesn't scale, without ML
- *"ML eats rules and code for lunch." – Elon Musk*

VISCOSITY

NYOUG

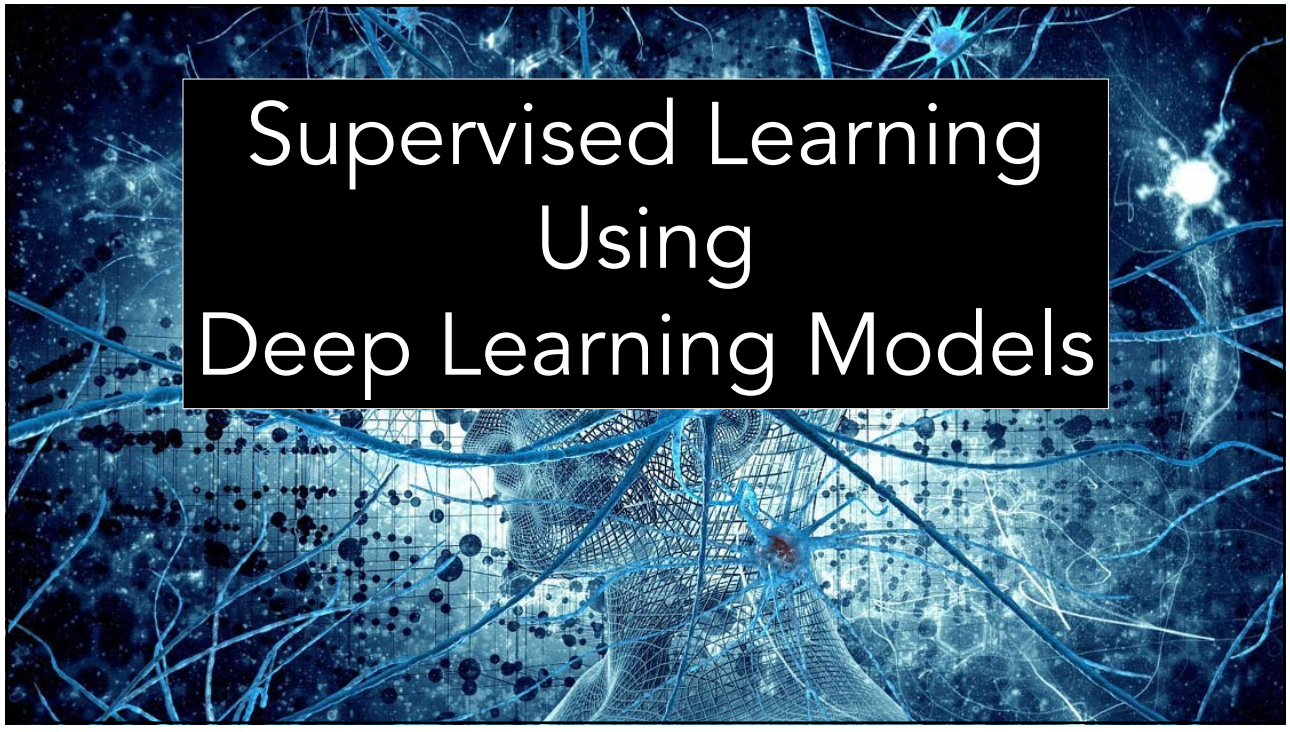

Artificial Neural Networks

- The history of artificial neural networks (ANN) began with Warren McCulloch and Walter Pitts in **1943**. I started playing with ANNs in 1995.
- ANNs are based on an understanding of **how our brains work**.
- A neuron **receives inputs**, sums them, if enough input results, the neuron **"fires" output** to its connected neurons.
- With enough input and neurons, then emergent "learning" behavior may occur.
- **How many?** Humans 86.00B, cats 0.76B, dogs 2.25B, sea squirt 231
- All the **learning is driven** by code and given objectives.
- We randomly disable neurons to make patterns more difficult to learn, called **regularization**.
- Neurons are typically **organized in layers**; 3 layers of 10 neurons each.
- Multiple layer **ANNs is nothing new** and not what Deep Learning is about.

VISCOSITY

NYOUG

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

How many neurons do creatures have?

Creatures	Neurons	Creatures	Neurons	Creatures	Neurons
African Elephant	257 000 000 000	Barn Owl	690 000 000	Cockroach	1 000 000
Human	86 000 000 000	Octopus	500 000 000	Honey Bee	960 000
Gorilla	34 400 000 000	Grey Squirrel	453 650 000	Fruit Fly	250 000
Brown Bear	9 586 000 000	Cockatiel	453 000 000	Ant	250 000
Dog	2 253 000 000	Pigeon	310 000 000	Jelly Fish	5 600
Racoon	2 122 000 000	Brown Rat	200 000 000	Sea Squirt	231
Cat	760 000 000	House Mouse	71 000 000	Sponge	0

dinoanimals.com/animals/number-of-neurons-in-the-brain-of-animals

VISCOSITY

NYOUG

Supervised Learning Artificial Neural Networks

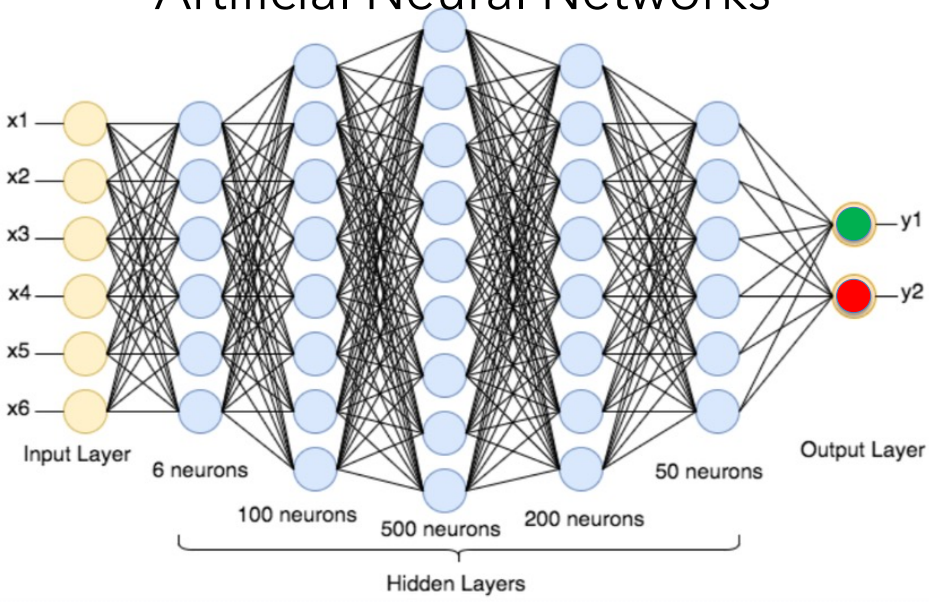

VISCOSITY

YOUG

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

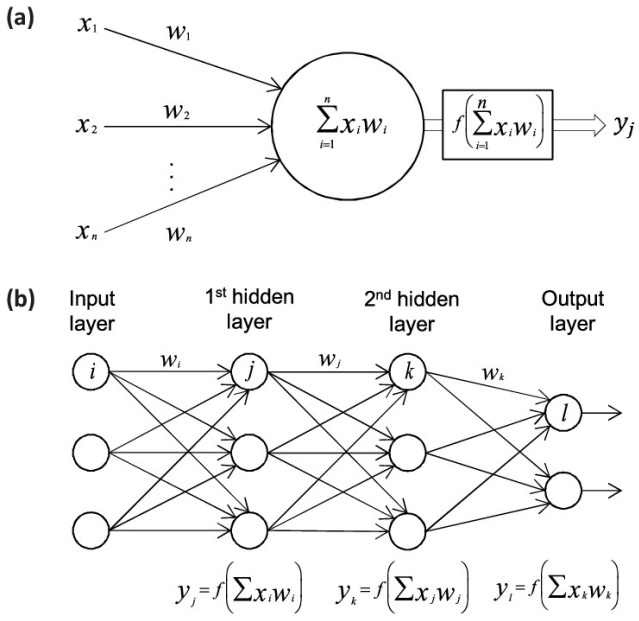

Supervised learning with Artificial Neural Networks

VISCOSITY

NYOUG

About Deep Learning #1

- **Deep Learning** models are "simply" more advanced Multilayer Percepton (MLP) Artificial Neural Networks (ANNs).
- Multiple layered ANNs have been around for many years.
- Neurons are now called "**cells**" because the cells are more complex and advanced than a classic artificial neuron.
- Each cell contains one or more "**gates**" and logic that is controlled through parameters and of course, weights.
- There are different highly **specialized** network **layer** types.
- The classic deep learning model is the **Recurrent Neural Network** or RNN for short. Cell **output is feed back** to its layer and also to the next neuron in its layer, providing sequence learning.
- **Keras** is a deep learning API written in Python, running on top of the ML platform TensorFlow. Keras allows a relatively simple and user-friendly way to define a neural networks.
- **TensorFlow** is an end-to-end open source platform for machine learning. It has a comprehensive, flexible ecosystem of tools, libraries and community resources that lets researchers push the state-of-the-art in ML and developers easily build and deploy ML powered applications. It was released in 2015 by Google.

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

About Deep Learning #2

- In many sequence applications,
  - Recent and older data are both important
  - A sample is related to its previous samples; **temporal data**
  - Different activation ("fire!) functions are needed
- RNNs have time-series sequence capabilities, but older patterns fade.
- LSTMs and GRUs have additional cell "gates" to improve older memory retention, plus other training capabilities.
- A number of (ever growing) deep learning architectures have been created to address this limitation.
  - Long Short-Term Memory – **LSTM**
  - Gated Recurrent Unit – **GRU**
- Convolutional Neural Networks (**CNN**) are used for images and video when combined with RNN type models.

Long Short-Term Memory NNs

- **Additional and specialized gates** within a LSTM cell, LSTMs have the ability to remember **sequences** giving them the ability to remember **patterns** of **both older and newer** data.
- They are very popular in speech recognition, music creation and time-series applications.
- LSTMs are highly flexible in structure. You can stack many hidden layers, combine with CNNs, used them in Encoder-Decoder networks, they can learn and predict in both time sequence directions, and generate new plausible sequences.
- **GRUs** provide these capabilities as well, but with **simpler cells**, which reduces training time.
- This is why I use LSTMs and GRUs when working on Oracle performance prediction projects.

VISCOSITY

NYOUG

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

Project

http://filebank.orapub.com/misc/Project-setup-6i.txt
http://filebank.orapub.com/misc/Project-AAS-REG-DL-v3j.ipynb

Build a real-time graphical prediction system using ML deep learning regression algorithm, multivariate dba_hist_sysmetric_summary data.

VISCOSITY

NYOUG

Project Details

- Supervised machine learning Deep Learning model is trained, new "real-time" data is polled and collected, prediction made, and the predicted and actual values are charted.
- The project uses many dba_hist_sysmetric_summary columns to build and create a real-time deep learning regression AAS prediction system that charts both the predictions and the actual AAS values.
- In this project we focus on the model building process steps, including model training and evaluation, preprocessing and prediction steps. The "real-time", always polling for new Sysmetric data, charting adds an interesting reality to this project.
- Google's Colab environment is being used: Jupyter like notebook, Virtual Linux server with "!" root access, GPU available, Google Drive R/W access, PC R/W access, and Internet access with "!" wget and scp.
- Python with all SciPy and Scikit-learn ML software libraries are available in Colab.
- Oracle Sysmetric data is used exclusively and provided by OraPub.

VISCOSITY

NYOUG

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on 24-May-2022 for the New York OUG. There is likely an updated version on the OraPub website at https://www.orapub.com

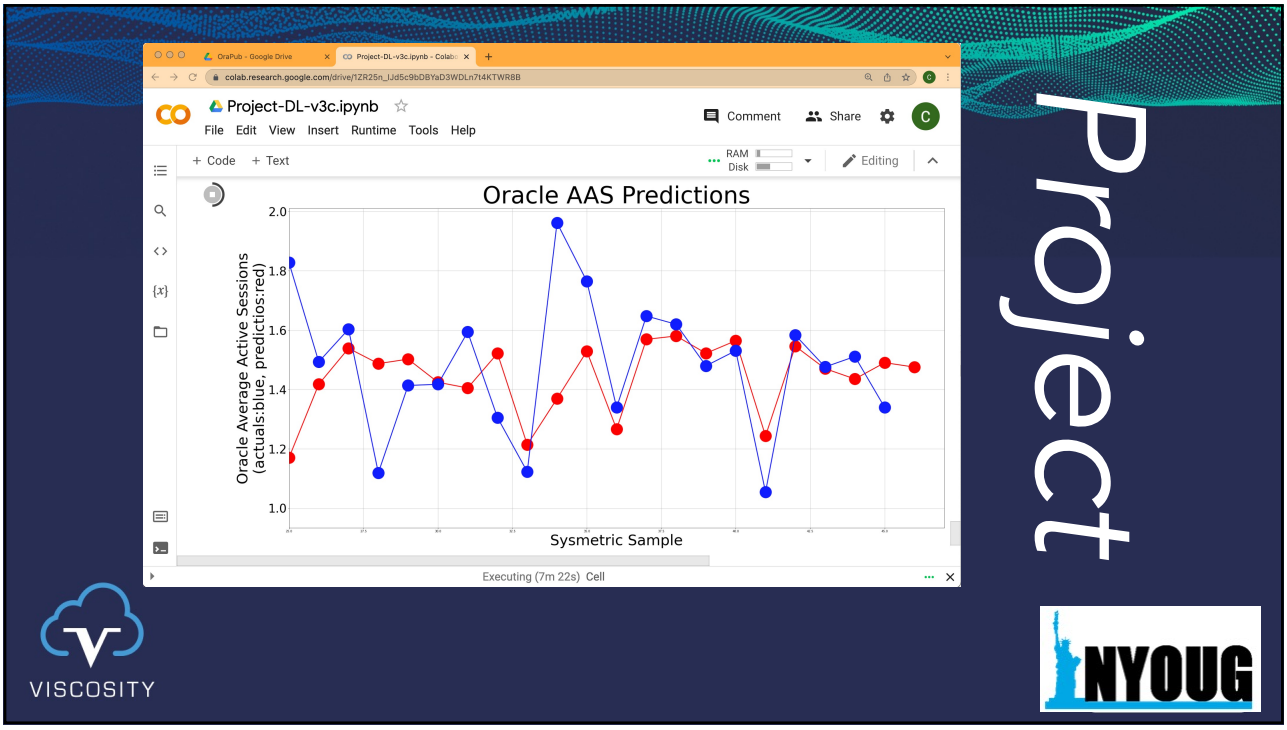


Going
Deeper


https://www.orapub.com/ml-ecourse
Machine Learning
For
Oracle Professionals
E-Course
ORAPUB
FREE from OraPub.com
OraPub is committed to cross-training
Oracle Professionals in Machine Learning

How to get **all the materials** for this presentation and updated versions later.

After the conference:
www.orapub.com/members > Tools & Presentations > Presentations

Resource listing

- **OraPub Membership** **for premium content**
  - "How To" Webinars – two each month. Over 130 recorded!!
  - Video Seminars – any device, any time, high quality
  - 20% LVC discounts!
  - Learning paths, mentoring, assessments and certificates, priority response
  - SLACK forum exclusively for paid members
- **Live Virtual Classroom (LVC) Training**
  - **ML 1 – Anomaly Detection & Deployment**
  - **ML 2 – Performance Prediction & Deployment**
  - Tuning Oracle Using An AWR Report
  - Tuning Oracle Using Active Session History (ASH) Strategies
- **Toolkits** – Many tools available at orapub.com
- **Craig's Blog & Website** – Search: "ml", "machine learning"
- **Presentations** – www.orapub.com
- **Books:** Oracle Performance Firefighting. Forecasting Oracle Performance.

VISCOSITY

NYOUG

Thank You!

VISCOSITY

NYOUG

AAS Predictions Using Deep Learning Neural Networks With Sysmetric Data

Craig Shallahamer
craig@orapub.com
craig@viscosityna.com

ORACLE ACE Director

VISCOSITY NORTH AMERICA

ORAPUB
POWERED BY VISCOSITY

NYOUG