DANIEL BUADES MARCOS

daniel@buad.es | +1 438-921-2325 | daniel.buad.es | dbuades | dbuades

SKILLS

Machine Learning & Data Science	**Pytorch**, **TensorFlow**, scikit-learn, pandas, Elastic Stack, Plotly Dash
Containers & Virtualization	**Kubernetes** (Helm, Kubeflow, Træfik, kubeadm), Docker, Linux KVM
DevOps	Github Actions CI/CD, IaC (**Terraform**, Ansible, Packer, cloud-init)
System Administration	Linux (Ubuntu Server, Arch), **Google Cloud**, AWS
Programming	**Python**, Bash, Java, SQL, LaTeX
Languages	English (114/120 TOEFL), French, Spanish

RELEVANT EXPERIENCE

Machine Learning Scientist Intern May 2018 - Jan 2020
R2 Inc. - Improving industrial processes safety with predictive maintenance. *Montreal*

- Prototyped and tested new deep-learning approaches to make real-time predictions on streaming data.
- Coded data ingestion pipelines, trained models distributedly and containerized deployments.
- Developed models based on Attention, LSTMs, CNNs, Autoencoders and other neural network architectures.
- Followed a CI/CD strategy to keep track of possible regressions and minimize concept drift.

Course Lecturer Sept 2018 – May 2019
Polytechnique Montréal - IND8217: Analytique de défauts et maintenance *Montreal*

- Statistics and machine learning applied to preventive maintenance and equipment failure prediction.
- Responsible for teaching and designing its contents, assignments and examinations.

EDUCATION

M.Sc.A. in Industrial Engineering - Machine Learning *Sept 2017 - Dec 2019*
Polytechnique Montréal
Canadian post-graduation work permit valid until 01/2023

B.Eng. in Industrial Technology Engineering
Universitat Politècnica de València, Spain *Sept 2013 - Jan 2017*
Polytechnique Montréal - *Recipient of an Erasmus+ KA107 grant* *Jan 2017 – Jun 2017*

PROJECTS

M. Thesis: A Deep Learning Approach for Condition-Based Fault Prediction in Industrial Equipment

- Developed a single deep learning model that replaced all the multiple expert-defined models used in production.
- Achieved 64% earlier fault detection, improving reliability while significantly reducing human intervention.
- NeurIPS 2019 LXAI Workshop - *Using a Self-Supervised Encoder for Predicting Faults in Electrochemical Cells*

Continous Sentiment Analysis on Twitter Streaming Data

- Collect tweets in real-time and index them in Elasticsearch for search, analysis and visualization with Kibana.
- Fine-tune different BERT-derived NLP models for the task and serve predictions over REST and gRPC APIs.
- Producer and consumers deployed as microservices on Kubernetes and decoupled via a RabbitMQ queue.

Self-Hosted Homelab

- Single Linux server hosting a virtual Kubernetes cluster composed of multiple KVM-based VMs and ZFS storage.
- Multiple microservices deployed in order to learn and experiment with new technologies.

Financial Statement Analysis of the ACS Group

- Perform fundamental and technical analysis of the sector and stock to make informed investment decisions.
- Research the company's assets, analyze their activities and search for correlations with competitors.