ChristosHadjinikolis

Lead ML Engineer @ Vortexa Ltd | Committee Expert Member @ JTC'21

christos.hadjinikolis@gmail.com

mediums

website: ML-AFFAIRS

LinkedIn

github

YouTube

languages

bilingual Greek/English

programming

Python

Java/Kotlin

SQL

Cypher

author

pypi:dynamicio

developer skills

TDD, git + CI/CD (Circleci)

cloud-platforms

GCP

AWS

Technologies & Frameworks

pySpark, Spark, Flink, Kafka, Elasticsearch, Neo4J, SQL, Docker, Kubelfow, Airflow

Python Libs

pandas, numpy

visualisations

Gephi, Seaborn, plotly

interests

Sports: Basketball, Football, Tennis, Jogging, Swimming

Reading: Comics, Philosophy, Politics

about me

Strategic ML Engineer with 14 years of combined academic and industry experience and a proven track record in transformative leadership and innovation. Championed key initiatives in MLOps, software architecture, and team mentorship, significantly enhancing operational efficiency and technical robustness. Recognized for spearheading event-driven development and stream-processing strategies, and driving cross-functional collaboration towards achieving organizational goals.

education

2010–2014 **Ph.D.** in Computer Science — King's College London, UK

***Thesis**: Persuasion Dialogues & Opponent Modelling.*
Developed a prediction mechanism applied on Knowledge Graphs. Involved the application of graph-analytics on large-scale graphs & statistical inference.

2004–2010 **Diploma (BEng)** in Computer Engineering — University of Thessaly, GR

A five-year-long polytechnic degree with a solid mathematical background (10 courses) ranging from pure (theoretical) computer science/engineering aspects to networking and telecommunications. *Majored in Artificial Intelligence*.

2002–2004 **Military Officer: Second Lieutenant** — National Guard, CY

Obligatory army service; position acquired through competitive IQ & athletic exams.

work-experience

since 12/2020 **Vortexa Ltd, London, UK** — Pod Lead | Lead ML-Engineer

Key member of the Data Production Team at Vortexa, I spearhead ETL/ML pipeline development for tanker operation forecasts. I specialise in transforming niche research into scalable real-world applications, enhancing fault-tolerance in processing pipelines. Collaborating closely with diverse experts, I bridge scientific experimentation with commercial viability and actively contribute to a decentralised conversational architecture practice in the R&D department as part of the Architecture Guild.

since 01/2021 **ISO - International Organisation for Standardisation** — Committee Expert Member for JTC'21

JTC'21 develops AI standards for European markets & society, aiming to align with EU policies & principles with international standards. Member of "Working Group 3: Engineering Aspects".

since 01/2015 **Kings College London, Dept. of Informatics, London, UK** — Associate Researcher

Interacting with members of the Agents & Intelligent Systems group on AI-related topics (my current interests orient around NLP applications in AI Logic, AI standardisation and AI ethics).

2016-2020 **DataReply (consulted for UBS & Vodafone), London, UK** — Lead ML Engineer

Transitioned from Data Scientist to ML Engineer, leading Data Science/Engineering teams. Responsibilities included translating business issues into analytics, overseeing project timelines, mentoring, & driving innovation. Experienced year-long roles at UBS, Vodafone & others.

2010–2016 **Roles in UK Institutions** — Associate Lecturer @GSM | Programming Teacher @DGC | TA @KCL|UCL

- **2012-2016**: Taught Computing, MATLAB, Java, Python, HTML, CSS at Greenwich School of Management (GSM) & David Game College (DGC).
- **2010-2014**: Teaching Assistant (TA) at King's College London (KCL) & UCL, covered Java Programming, Data Structures, Computer Systems, Semantic Web, DBMS-SQL, & AI.

Consulting Roles with Data Reply

Vodafone Group London [09/2019 - 11/2020] Senior / Lead ML Engineer

Key Contributions & Projects:

- **ML Framework Development & Migration Supervision:** Led the migration of ML models and CI/CD pipelines to GCP, focusing on pySpark-based ML industrialization and development of base images for DataProc and TensorFlow.
- **Developing Cloud-native Data Analytics on GCP:** Served as technical lead for cloud-native data analytics projects, specialising in Big Data solutions for marketing, recommendations, automation, and resource optimisation.
- **Infinity Platform:** Drove the implementation of CI/CD pipelines within an Apache Kubeflow-based platform, enhancing the development of ML pipelines in the Advanced Analytics Team.

Technologies: *Python, pySpark, Jenkins, AirFlow, Kubernetes, Docker, GCP, KubeFlow*

CNHi, Chicago [05-09/2019] Head of Data Science Role

Analysis of Agriculture Vehicle Data

Description: Led a data science and engineering team, translating business requirements into actionable tasks. Focused on statistical analysis and developing time series prediction models for threshold-based alerts.

Technologies: *Azure*, *Databricks*, *PowerBi*, *PySpark*, *Python (sklearn, pandas)*

UBS, London [08/2017-04/2019] ML Engineer Role

Graph Analytics & Data Processing Engineer

Description: Developed end-to-end processing pipelines to ingest and analyse data in real-time. Applied Graph Analytics for the extraction of actionable insights, challenging organisation structures and aiding process optimisation (batch analysis). Applied Process Mining techniques to aid understanding of various processes and behaviours of entities of interest in event logs (stream processing). Generated useful visuals through dashboards and Python notebooks.

Technologies: *Elasticsearch*, *Kafka*, *Apache Flink*, *Python (sklearn)*, *Kibana*

Irish Road Safety Authority, London [05-07/2017] Data Science Role

Analysis of Road Safety Data

Description: Involved analysing and extracting insights from MySQL data for predicting highway crashes, identifying high-risk nodes in the traffic network as well as categorising/profiling drivers/operators based on compliance rates, inspections and defects.

Technologies: *MySQL*, *Neo4J*, *Google Cloud*, *R*, *Python (sklearn)*

GameSys, London [03-04/2017] Data Engineering Role

Architecture Design and Support for new Digital Platform

Description: POC work concerned with showcasing how a Kafka Schema Registry can be used for data validation purposes. Trained a team of Engineers on how to deploy it and integrate it into their existing solution.

Technologies: *Apache Kafka*, *Google Cloud*

The Mondadori Group, Milan [11–12/2016] Data Science/Engineering Role

Development of a Recommender Engine

Description: Analysed user preferences harvested from this publisher's social commerce platform where members can interact by sharing/buying/rating books maintained in their digital "library". Applied community detection & market basket analysis techniques (ALS method) and the extraction of book association rules, all of which could then be used for both user-to-user and user-to-book recommendations.

Technologies: *Scala, SQL, Apache Spark, MongoDB*

Unicredit, Milan [10–11/2016] Data Engineering Role

Development of a Fuzzy Matching Mapping Engine

Description: Developed a batch process solution for entity identification and matching between data retrieved from distinct databases, based on a variant of Levenshtein's distance algorithm.

Technologies: *Scala, SQL, Apache Spark*

Data-Science/Engineering Projects

Python Library Author [01-04/2022] Engineering

Main Description: Dynamicio: A pandas wrapper for abstracting I/O code

Dynamicio (or dynamic(i/o)) is a convenient wrapper around pandas I/O operations. It's a manifestation of the dependency inversion principle–a layer of indirection if you want–which keeps your code DRY and increases re-usability, effectively decoupling business logic from the I/O layer. It supports:

- seamless transition between environments;
- abstracting away from resource and data types through resource definitions;
- honouring your expectations on data through schema definitions;
- metrics auto-generation (logging) for monitoring purposes.

Github Repo: here

Side Project [01-03/2018] Data Science Side Project

Prototyping: Feature-based Customer Review Mining- An Implementation

Description: This project, summarises and explains the approach elaborated in [Kushal Bafna & Durga Toshniwal, 2013]. Its purpose was to serve as a prototype for an internal company project I was working on. The problem I addressed concerns allowing customers to extract objective review summaries relevant to a book, without having to go through the trouble of reading all of the available reviews.

Github Repo: here

Technologies: *Python (nltk), Association Rules Mining, Python Notebooks*

DataReply UK, London [01-03/2017] Data Science Role, Internal Project

Team R&D Project - Conference Demonstration: Building a chatbot

Description: Used a corpus of movie dialogues to train a sequence-to-sequence model for the implementation of a conversational chatbot.

Technologies: *Python (nltk), Apache Spark, s2s NNs, Tensorflow, Google Cloud*

DataReply UK, London [09–10/2016] Data Science/Engineering Role, Internal Project

Team Project - Conference Demonstration: Facial expression recognition application

Description: Developed a Neural Network model capable of recognising facial expressions in images; then replaced faces in an image with the corresponding emoticon. The model had to be fed with a large number of categorised images. Image collection was achieved through web-scrapping (beautiful-soup) + semi-manual filtering.

Technologies: *Python, OpenCV, Keras, Tensorflow*

DataReply UK, London [05–09 2016] Data Science/Engineering Role, Internal Project

R&D Project: Graph analytics & Data Augmentation

Description: Developed a graph analytics solution for company-director networks. Local analysis: developed a real-time graph neighbourhood constructor for companies of interest with multiple visualisation options. Global analysis: carried out centrality (GraphFrames & PageRank) and community detection analytics.

Github Repo: here

Technologies: *Scala, Apache Spark, CYPHER, Neo4j, Python, webAPI*

DataReply UK, London [06 2016] Data Science/Engineering Role, Internal Project

Conference Demonstration Project: Twitter Sentiment Analysis

Description: Used the twitter API to collect MP tweets related to #BREXIT. Tweets were then categorised based on the public position of the respective MP and were analysed using NLP libraries. Finally, a naive Bayes prediction model was developed to predict any user's opinion, by leveraging their own tweets and those of people they follow.

Github Repo: here

Technologies: *Twitter API, Python (nltk, pySpark), Apache Spark*

certificates & awards

07/2020	**AWS Certified Machine Learning - Specialty** Validation Number: D0R0WC7CN1VQ13SV	AWS
01/2019	**Process Mining** License number: AFU2VTP6273V	Coursera Course Certificates
10/2017	**Google Certified Professional - Data Engineer** Credential ID: 000000400	Google Cloud Certified
07/2017	**Advanced Elasticsearch: Data modelling** Enrolment Id: 33442	Elastic
06/2016	**Graph Analytics for Big Data** License number: NZADYAQ7N8FW	Coursera Course Certificates
05/2016	**Apache Spark Developer Certification** License number: 1.x–0548	O'Reilly Media
04/2016	**Neo4j Certified Professional** License number: here	Neo Technology, Creators of Neo4j
2013	**Best Poster Award** Selected amongst 413 contestants (Prestigious Conference)	IJCAI Conference, Beijing, China
2012	**Outstanding Teaching Assistant Award** Notable Mention: Department of Informatics (based on student voting)	King's College London
2011	**Graduate Certificate in Academic Practice** Related with improving one's teaching abilities	King's College London
2010-2013	**PhD Tuition Fees Scholarship** Provided by EPSRC and King's College London	EPSRC

talks/interviews

2023	**Bill Raymonds, Agile in Action** **Agile Data Science**: Sailing Through Data Science: The Agile Journey at Vortexa, **Details:** here	Podcast Interview
2022	**Open Data Science Conference** **Dynamicio**, a python library for abstracting your I/O, **Details:** here	Industry Conference Talk
2020	**iunera: Big Data Innovators** The Agile approach in data science explained by an ML expert, **Details:** here	Interview Blog Post
2020	**Big Data Warsaw** Monitoring & Analysing Communication and Trade Events as Graphs, **Details:** here	Industry Conference Talk
2018	**Connected Data London** AI tribes and interconnections: what's graph got to do with it?, **Details:** here	Industry Conference Talk
2018	**Minds Mastering Machines** A talk on doing Data Science the Agile Way, **Details:** here	Industry Conference Talk
2018	**London Big Data Week** A talk on Big Data and Graph Analytics, **Details:** here	Industry Conference Talk
2016	**Open Data Science Conference** A talk on the advancements in the field of graph-isomorphism, **Video:** here	Industry Conference Talk
2013	**The Imperial College London ACM Student Chapter Seminars** On anticipating opponent information in competitive contexts, **Video:** here (Invited)	Workshop Talk
2011	**The 6th Joint Working Group Workshops, Maastricht University, NL** On Strategic Consideration in Dialogue Games (Invited)	Workshop Talk

Academic Achievements: Publications & Committees

2016 **MIWAI**
PC Member
The 10th Multi-Disciplinary International Workshop on AI

2015 **MIWAI**
PC Member
The 9th Multi-Disciplinary International Workshop on AI

2015 **Building Support-based Opponent Models in Persuasion Dialogues**
C. Hadjinikolis, S. Modgil, E. Black
Theory & Applications of Formal Argumentation (LNAI) (Workshop)

2013 **Opponent Modelling in Persuasion Dialogues**
C. Hadjinikolis, Y. Siantos, S. Modgil, E. Black, P. McBurney
23^{rd} International Joint Conference on AI (AAAI)(Highly recognised)

2012 **Mechanisms for Opponent Modelling**
C. Hadjinikolis, S. Modgil, E. Black, P. McBurney
Imperial College Computing Student Workshop (*Schloss Dagstuhl*)(Workshop)

2012 **Investigating Strategic Considerations in Persuasion Dialogue Games**
C. Hadjinikolis, S. Modgil, E. Black, P. McBurney, M. Luck
Starting AI Researcher Symposium (IOS Press)(Workshop)

2012 **Minimising Processor Communication in Parallel Approximate String Matching**
C. Hadjinikolis, C. S. Iliopoulos, S. P. Pissis, and A. Stamatakis
Heidelberg Institute for Theoretical Studies (Tech. Report)

2012 **ICCSW**
PC Member
Imperial College Computing Student Workshop

2012 **STAIRS**
PC Member
6th European Starting AI Researcher Symposium