

CV // JOÃO GALEGO

[in linkedin.com/in/jgalego](https://www.linkedin.com/in/jgalego) | github.com/JGalego | jgalego.github.io

OVERVIEW

João Galego is a Principal AI Engineer at Critical Software and an Invited Professor at ISEG. He has a MSc in Physics from IST and a post-graduate degree in Forensic Studies from ISCSP. Before Critical Software, he worked as a DevOps Engineer and Lead ML Engineer at Siemens and as a Solutions Architect, ML Specialist, and Quantum Ambassador at AWS, helping startups at every stage of their cloud journey. When he is not working or teaching, you can usually find him browsing the local bookstore, contributing to open source projects, or enjoying some quality family time.

PROFESSIONAL EXPERIENCE

| | |
|--|----------------|
| Principal AI Engineer Critical Software | 2025 – Present |
| Building dependable AI solutions for a better and safer world. | |
| Startup Solutions Architect AWS | 2022 – 2025 |
| Helping startups at every stage of their cloud journey. | |
| DevOps Engineer / Lead ML Engineer Siemens | 2019 – 2022 |
| Managing industrial-grade applications in the cloud. | |
| DevOps Engineer Innowave | 2017 – 2019 |
| Moving faster from idea to production. | |
| Test Automation Engineer Coriant | 2014 – 2017 |
| Scaling software testing through automation. | |

EDUCATION

| | |
|--|-------------|
| PhD Cognitive Science (ABD) Universidade de Lisboa | 2017 – 2020 |
| Studied motor cognition by bringing minds and computers together using ML. | |
| PgD Forensic Studies 🎓 ISCSP | 2016 – 2017 |
| Learned what "every contact leaves a trace" actually means. | |
| MSc Physics 🎓 IST | 2008 – 2013 |
| Failed to solve the riddle of baryon asymmetry. | |

TEACHING

Applied Artificial Intelligence & Machine Learning

ISEG

2024 – Present

Artificial Intelligence for Value Creation

ISEG

2024 – Present

CERTIFICATIONS



Oxford Machine Learning Summer School

AI for Global Goals

2024

AWS Certified Solutions Architect - Professional

AWS

2024

AWS Certified Database - Specialty

AWS

2023

Oxford Machine Learning Summer School

AI for Global Goals

2022

Startup School

Y Combinator

2022

AWS Certified Solutions Architect - Associate

AWS

2022

Migration Ambassador Foundations (Business)

AWS

2022

AWS Partner: Accreditation (Technical)

AWS

2022

AWS Partner: Cloud Economics Accreditation

AWS

2022

Machine Learning Engineering for Production (MLOps) Specialization

DeepLearning.AI

2021

Natural Language Processing (NLP) Specialization

DeepLearning.AI

2021

Google Data Analytics Professional Certificate

Google

2021

AI Product Manager Nanodegree

Udacity

2021

Self-Driving Car Engineer Nanodegree

Udacity

2021

AWS Certified Machine Learning - Specialty

AWS

2020

Fundamentals of Digital Marketing

Google

2020

Microsoft Certified Azure Fundamentals (AZ-900)

Microsoft

2020

| | |
|--|------|
| Associate Android Developer Google | 2018 |
| Certified Jenkins Engineer Cloudbees | 2016 |
| Machine Learning Coursera | 2014 |

TECHNICAL SKILLS

Disclaimer: *This list is neither extensive nor exhaustive by design. As a firm believer in the “right tool for the right job” mindset, I’ve had the chance to work with hundreds of tools, packages and frameworks throughout my career. Listing them all here would be cumbersome and, worse, uninformative. Please contact me directly if you wish to discuss specific technical skills.*

- Hands-on experience with multiple cloud providers (AWS, Azure, Google Cloud).
- Active contributor to several open source projects (AWS, Google DeepMind, Hugging Face, LangChain, Stanford NLP)
- Working knowledge of multiple programming and scripting languages (C, C++, Python, Java, Rust)
- Expert knowledge of multiple DevOps tools (Ansible, Terraform, GitLab CI, Docker, Kubernetes)
- Extensive experience with common Python packages for scientific computing (NumPy, SciPy), data analysis (Pandas, Vaex) and differentiable programming (JAX, MXNet, PyTorch, TensorFlow), NLP modules (Gensim, NLTK, spaCy, TextBlob), and high-level ML frameworks (HuggingFace Transformers and Diffusers, Keras, Scikit-Learn, Trax).
- Expert knowledge of data and ML engineering tools/frameworks

HONORS & AWARDS

| | |
|--|------|
| EMEA CS SA GenAI Builders Quest - 1st Place, GenAI Trailblazers AWS | 2023 |
| IBM Quantum Challenge - Advanced, 1st Place IBM | 2023 |
| Hugging Face 🤗 Deep Reinforcement Learning Course - Certificate of Excellence Hugging Face | 2023 |
| IBM Quantum Challenge - Advanced, 1st Place IBM | 2022 |
| Machine Learning University - AutoGluon - Top 10 AWS | 2022 |
| Machine Learning University - GML - Top 10 AWS | 2022 |
| QHack 2022 QML Challenge - Top 70 Xanadu | 2022 |
| IBM QGSS21 Certificate of Quantum Excellence IBM | 2021 |
| Google Cloud Hero: BigQuery - Analytics & Machine Learning - 1st Place Google / Siemens | 2021 |
| IBM Quantum Challenge - Advanced, 1st Place IBM | 2021 |

| | |
|---|------|
| Udacity Technology Scholarship - AI Track Udacity / Bertelsmann | 2021 |
| QHack 2021 QML Challenge - Top 50 Xanadu / AWS | 2021 |
| SITEX 2021 - Industrial Edge, 1st Place Siemens | 2021 |
| AI@Sustainability Hackathon - Technical Excellence (Team Mobility) Siemens | 2020 |
| Associate Android Developer Fast Track Program Scholarship Google / Udacity | 2016 |
| Junior Research Scholarship Fundação para a Ciência e Tecnologia (FCT) | 2013 |

ARTICLES & BLOGS

| | |
|--|------|
| Hacking GraphRAG with Amazon Bedrock AWS Community | 2024 |
| Learn how to run GraphRAG pipelines backed by Amazon Bedrock using LiteLLM proxy. | |
| Make Programs, not Prompts: DSPy Pipelines with Llama 3 on Amazon SageMaker JumpStart AWS Community | 2024 |
| Learn how to create ML pipelines with DSPy powered by Meta's Llama 3 70B Instruct model running on Amazon SageMaker. | |
| Leaving no language behind with Amazon SageMaker Serverless Inference AWS Community | 2024 |
| Host a high-quality translation model at scale using Amazon SageMaker Serverless Inference. | |
| Build Document Processing Pipelines with Project Lakechain AWS Community | 2024 |
| Learn how to create cloud-native, AI-powered document processing pipelines on AWS with Project Lakechain. | |
| Your idea is my command: multi-agent collaboration for software development AWS Community | 2024 |
| Create your own virtual software company with Amazon Bedrock and ChatDev and use it to build custom applications through LLM-powered multi-agent interactions. | |
| Running LangChain.js Applications on AWS Lambda AWS Community | 2024 |
| Learn how to run LangChain.js apps powered by Amazon Bedrock on AWS Lambda using function URLs and response streaming. | |
| Scrape All Things: AI-powered web scraping with ScrapeGraphAI and Amazon Bedrock AWS Community | 2024 |
| Learn how to extract information from documents and websites using natural language prompts. | |
| Bug hunting with Amazon Bedrock and SWE-Agent AWS Community | 2024 |
| Use Amazon Bedrock with SWE-Agent to create your own software engineering agent that can fix real-life bugs and issues in GitHub repositories. | |

- Fighting Hallucinations with LLM Cascades** 2024
AWS Community
Learn how to implement FrugalGPT-style LLM cascades on top of Amazon Bedrock using LangChain and LangGraph without breaking the bank.
- Running Small Language Models on AWS Lambda** 2024
AWS Community
Learn how to run small language models (SLMs) at scale on AWS Lambda with function URLs and response streaming.
- Mapping embeddings: from meaning to vectors and back** 2024
AWS Community
My adventures with RAGmap and RAGxplorer with a light introduction to embeddings, vector databases, dimensionality reduction and advanced retrieval mechanisms.
- Building GenAI apps with managed AI services** 2024
AWS Community
Some random thoughts on managed AI services and their place in the GenAI stack... with examples.
- The Rise of the LLM OS: From AIOS to MemGPT and beyond** 2024
AWS Community
A personal tale of experimentation involving LLMs and operating systems with some thoughts on how they might work together in a not-so-distant future.
- YOLambda: Running Serverless YOLOv8/9** 2024
AWS Community
Run inference at scale with YOLOv8/9 in a secure and reliable way with AWS Lambda.
- Deploy LangChain applications on AWS with LangServe** 2024
AWS Community
Learn how to deploy LangChain applications with LangServe in minutes on Amazon ECS and AWS Fargate using AWS Copilot.

TALKS & EVENTS

- O novo paradigma de GenAI com Amazon Bedrock** 2024
AWS Generative AI Day
L200 talk on Amazon Bedrock and GenAI for everyone.
- Demistifying ML with Amazon SageMaker: Agility, Speed and Efficiency** 2023
AWS Cloud Experience Portugal
L200 talk on low-code and no-code ML with Amazon SageMaker.
- How to solve common business challenges with AI/ML** 2022
AWS Initiate Portugal 2022
L200 talk featuring Sérgio Casaleiro (Software Development Manager at Ryanair Labs) on how to leverage the AWS ML Stack to solve common business problems and move an AI-powered application from PoC to production in no time.

Amazon SageMaker for Startups

2022

AWS Experience Madrid on Floor 11

L300 talk on how to use Amazon SageMaker to manage the entire MLOps pipeline at scale - from data exploration and feature engineering to model deployment and monitoring - with a focus on cost optimization.

Machine Learning for Medicine and Healthcare

2019

12th Workshop on Biomedical Engineering (WBME)

An overview of what ML actually is, highlighting some of its applications to the health sector - from the rise of expert systems in the 80s to the diagnosis, prognosis and treatment of COVID-19 - how it is shaping the present and how it may one day decide our future. Slides are available [here](#)