

Type d'offre : Corporate offer

Post date : 03.06.25

DeepFlows

CTO Late Co-Founder

Informations générales

Contract type : Permanent contract

Contact :

[Greg Batby](#)

Starting date : Mon 01/09/2025 - 12:00

DeepFlows :

Founded in 2024 by ML & Soft engineers from Google, Polytechnique, Centrale Paris, and a former investment banker, [DeepFlows](#) is a fast-growing technology company that automates tedious analytical work for founders, financial advisors and investors. [Backed by private investors from J.P. Morgan, Morgan Stanley, Meta, and others](#), and [backed by Microsoft](#), the platform leverages cutting-edge AI agents to radically accelerate and improve decision-making in M&A, private equity and asset management.

Détail de l'offre (poste, mission, profil) :

Key Responsibilities & Tech Stack

Domain	Key Missions
--------	--------------

<i>AI Pipeline Dev.</i>	Design, build, and optimise a multi-agent orchestration layer (CrewAI → LangChain), plug-and-play LLM back-ends (Azure OpenAI, Amazon Bedrock), Retrieval-Augmented Generation pipelines (vector + knowledge-graph RAG in Neo4j), and a full evaluation/observability loop
<i>MLOps & Infra</i>	Own end-to-end CI/CD for models and services; containerise with Docker, run at scale on AKS/EKS, and govern experiments via Azure ML + MLflow. Stand-up AWS pathways for Bedrock-native models and future hybrid workloads
<i>Data & Knowledge Engineering</i>	Ingest and transform high-volume financial data using PySpark and Pandas; power low-latency search with FAISS vectors and Azure Cognitive Search; manage graph relationships in Neo4j and analytics in Azure Synapse
<i>Security & Compliance</i>	Unqualified SOC 2 Type II audit; zero major pen-test findings
<i>AI Product Strategy</i>	Translate fundraising and investor workflows into NLP/LLM solutions; drive model-selection, prompt-tuning, RAG/KG innovation, and cost-to-accuracy optimisation; maintain a forward-looking tech radar for emerging foundation models
<i>Leadership & Culture</i>	Set engineering standards, run design reviews, and foster knowledge-sharing through Linear, GitHub, Azure DevOps, and living internal docs; mentor a high-performance team and recruit top-tier AI talent

Qualification & Required Skills

- Degree in computer science, software engineering, or related field, or equivalent experience;
- At least +4-6 years of professional experience in Machine Learning Engineering and deep hands-on expertise with LLMs, retrieval-augmented generation, and knowledge graphs;

- Proactive and capable of working independently, yet thrive in collaborative team environments;
- Comfortable interfacing with customers, investors, and the board;
- **Bonus:** Published research or open-source contribution in NLP . Experience with Finance/Investment firms. Prior success in a seed to series B startup journey.

URL de l'offre :

<https://www.dataia.eu/sites/default/files/2025%20-%20Deepflows%20-%20CTO.pdf>

Lien vers l'offre sur le site dataia.eu :<https://da-cor-dev.peppercube.org/node/1311>