




CSO / Late Co-Founder – Equity & Salary Package – Miami/Dubai/Remote

 Founded in 2024 by ML & Soft. Engineers from Google, Polytechnique, Centrale Paris, and a former Investment Banker, DeepFlows is a fast-growing tech company automating time-consuming analytical work for founders, financial advisors and investors.

 Backed by Private Investors from [J.P. Morgan](#), [Morgan Stanley](#), [Meta](#), and others, and supported by [Microsoft](#) our platform leverages cutting-edge AI agents to drastically accelerate and improve decision-making in M&A, Private Equity, and Asset Management

 We're working hand-in-hand with Founders and Investment Professionals to build a verticalized multi-AI-Agents platform that supports the drafting and review of memos, due diligence reports, market analysis, and much more. We now aim at becoming an All-in-one application for Founders, Financial Advisors and Investors.

Role Description

As **CSO**, you will lead DeepFlows' scientific and research direction, with a focus on pushing the frontiers of applied AI in financial workflows. You'll drive long-term AI/ML research, define our AI strategy in coordination with product and engineering, and ensure that our agents stay ahead of the curve in terms of accuracy, innovation, and performance.

You'll collaborate closely with the CEO and CTO, serve as the primary architect for model innovation (RAG, LLMs, knowledge graphs), and represent DeepFlows' research in academic, industry, and investor circles. Your work will directly inform the models deployed in production and shape how we approach reasoning, search, evaluation, and trust in multi-agent environments.

Key Responsibilities & Tech Stack

Domain	Key missions
AI Research Strategy	Define the scientific roadmap for DeepFlows' agents; align model evolution with product goals and market trends
Multi-Agent Reasoning	Research and improve orchestration, coordination, and memory across LLM-based agents
Retrieval-Augmented Generation	Design and improve RAG pipelines (vector + graph) for financial document analysis
Evaluation & Alignment	Build internal benchmarks and observability tools to evaluate hallucination, latency, cost, and decision confidence
Model & Prompt Innovation	Own experimentation across model families (GPT-4, Claude, Mistral, etc.); drive prompt engineering and adaptive agent behavior
Research & Publication	Contribute to the AI community through whitepapers, open-source work, or conference publications (NeurIPS, ACL, etc.)

Qualification & Required skills

- PhD (or equivalent) in Machine Learning, NLP, Mathematics, or related fields
- 4–6+ years of experience in AI/ML research or applied LLM research
- Deep understanding of retrieval-augmented generation, LLM evaluation, and scientific modeling
- Track record of delivering cutting-edge innovation in real-world products or research environments
- Bonus: experience with financial data, investment workflows, or working with structured + unstructured documents
- Bonus: publications, patents, or open-source contributions in AI/NLP/IR