

Type d'offre : Offre en entreprise

Date de publication : 03.03.26

Formel AI

AI Engineer

Informations générales

Type de contrat : CDI

Durée du contrat : CDI

Niveau d'expérience : 3 years and more.

Contact :

careers@formel.ai

Date de prise de poste : mer 01/04/2026 - 12:00

Formel AI :

At [Formel AI](<https://formel.ai/>), we combine the creativity of generative AI with the reliability of mathematical formal methods. By leveraging the verification capabilities of formal methods, such as the Lean 4, we pioneer a new generation of Large Language Models (LLMs) that are reliable, transparent, and cost-efficient. We're honored to have received the [2nd spot of the Innovation Prize from École normale supérieure](https://www.linkedin.com/posts/ecolenormalesuperieure_jeudi-soiront-%C3%A9t%C3%A9-annonc%C3%A9s-les-premiers-ugcPost-7401284253482635264-odbQ) in Nov 2025.

Formel AI is founded by [Sylvain Combettes](<https://www.linkedin.com/in/sylvain-combettes/>) (CEO) and [Antoine Mazarguil](<https://www.linkedin.com/in/antoine-mazarguil/>) (CTO), who met during their PhDs at ENS Paris-Saclay. We closed an [oversubscribed angel round](https://www.linkedin.com/posts/sylvain-combettes_with-antoine-mazarguil-were-thrilled-activity-7429435501692006400-i3u0) in Feb 2026, backed by 30+ top-tier business angels including the founders of Datadog, Dataiku, Nabla, and Artefact. Our scientific advisory board includes Full Professors from Ecole polytechnique, ENS Paris-Saclay, and Ecole des Ponts.

We're now onboarding our first design partners and building the team to make it happen. Come join us!

Adresse :

Formel AI
75000 Paris
France

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

About the role

As our **AI Engineer**, you will be at the forefront of AI innovation, working on groundbreaking language models with formal verification methods. You'll join a founding team of deep-tech engineers tackling some of the most challenging problems in AI: building LLMs that are transparent, reliable, and cost-effective.

This is a hands-on research and engineering role where you'll design and develop the next generation of AI systems. You'll work on novel architectures, orchestration protocols, and training methodologies that dramatically reduce computational costs.

You'll be joining at a critical inflection point—moving from research breakthroughs to real-world, high-stakes domains like healthcare and enterprise software, where AI reliability is non-negotiable.

Key responsibilities

- **Design and develop AI systems:** Lead the architecture and implementation of LLMs integrated with our existing systems
- **Build orchestration protocols:** Design and implement orchestration frameworks, including training, deployment, and verification workflows
- **Drive research innovation:** Actively shape the company's research roadmap, exploring new methodologies, and verification techniques in collaboration with our scientific advisory board
- **Deploy with design partners:** Work closely with 2-3 early customers in high-stakes domains to build solutions that demonstrate measurable improvements in AI reliability and cost-efficiency
- **Optimize training and evaluation:** Develop advanced dataset generation pipelines, fine-tuning, and rigorous evaluation protocols to continuously improve model performance
- **Build reusable infrastructure:** Create documentation, code templates, and engineering best practices to scale our deployment approach as we grow
- **Bridge theory and practice:** Translate cutting-edge research in AI into production-grade solutions

What we are looking for

Essential qualities

- **Research-driven curiosity:** You're passionate about exploring novel AI applications and are curious about what's possible
- **Hands-on technical excellence:** You can design, implement, and deploy production-grade AI systems from model training to orchestration
- **Deep technical expertise:** You have deep knowledge of LLM architectures, training methods, and evaluation techniques
- **Bridge builder:** You can translate complex research concepts into practical implementations

Background

- 3+ years of hands-on experience building and deploying AI/ML systems, with deep expertise in cloud infrastructure
- Strong foundation in machine learning fundamentals, model architectures, and modern AI engineering
- Deep knowledge of the latest developments in AI research and production systems

Required skills

- **Python mastery:** You write clean, efficient Python code and understand object-oriented programming
- **Fullstack fundamentals:** You understand both frontend and backend development principles and are proficient in Python for backend systems

- **LLM expertise:** You have hands-on experience with large language models, including API in orchestration frameworks
- **LLM inference packages:** You have practical experience with at least one major LLM inference framework (vLLM, TensorRT-LLM, or similar) and understand inference optimization techniques
- **Training & fine-tuning:** You've implemented fine-tuning workflows (DPO, RLHF, or similar) and model evaluation
- **Model architectures:** You have deep knowledge of transformer architectures and modern
- **Production experience:** You've deployed AI systems in production environments and understand the full lifecycle from training to deployment
- **Research mindset:** You stay current with the latest AI research and can quickly prototype

Benefits

- Financial & perks

- **Competitive cash salary** and equity (BSPCE) in a high-potential company.
- **Health insurance** (Alan) for you and your family (partner and kids).
- **Daily lunch** vouchers (Swile).
- **Hybrid work** model with a strong emphasis on in-office collaboration (typically 3-4 days per week) to ensure team cohesion, spontaneous discussions, and deep work together.

- What makes this role unique:

- **Founding team impact:** Join as one of the first employees and shape the trajectory of a high-growth startup.
- **Mission-driven startup:** Help solve the key pressing challenges of AI—the most transformative technology of our time. Our LLMs are trustworthy, transparent, and cost-effective, with direct benefits for environmental sustainability.
- **Work with cutting-edge technology:** Build systems that combine LLMs with formal verification, operating at the forefront of AI safety and reliability where breakthrough developments happen every week.
- **World-class team:** Collaborate with a scientific advisory board from leading institutions (MIT, Stanford, etc) and experienced business angels from top tech companies (Datadog, Dataiku, Nabla, Arteris, etc).
- **Strategic customers:** Work with leading companies in healthcare, enterprise software, and fintech.
- **Flexibility of an early-stage startup:** Autonomy to define how you work and what you work on.

Hiring process

1. Apply and submit your resume using this [Google Form](<https://docs.google.com/forms/d/e/1FAIpQLSfvLZeeEiMIhtFWlmd2lbaQ9ghPhYw3r0nSR>)

2. **Screening call (30 min):** A quick chat with the CTO to align on vision and expectations.
3. **In-depth technical discussion** (home exercise followed by a 60 min discussion). Best way to decide if we are a great fit as potential collaborators.
4. **Technical deep-dive (60 min):** A focus on your proficiency with the tools you will use.
5. **Founder interview (45 min)** Final discussion on strategy, culture fit, and your long-term evolution within the company, with the CTO and CEO.

URL de l'offre :

<https://docs.google.com/forms/d/e/1FAIpQLSfvLZeeEiMIhtFWlmd2IbaQ9ghPhYw3r0nSRVHZejb>

Lien vers l'offre sur le site dataia.eu :<https://da-cor-dev.peppercube.org/node/1535>