



Application for Junior ML Engineer

About the job:

Location: Hyderabad (onsite)

Experience: 0–2 year (including internships / strong projects)

About the role

We are building an internal AI system that can read complex technical/scientific documents, extract structured SOP-style procedures, learn from them over time, and power predictive tools for decision-making. As a Junior ML Engineer, you will work end-to-end on this system using open-source, self-hosted LLMs running on our own RTX-class hardware, not public cloud LLM APIs.

This role is ideal for someone who is strong in Python, genuinely excited about LLMs and RAG, and wants very high ownership in a small, early-stage setup.

Responsibilities

- Build and maintain pipelines that parse PDFs and other documents into clean text, then convert them into structured SOP-style outputs
- Use self-hosted, open-source LLMs together with RAG to extract information and answer questions over our internal document and knowledge bases.
- Design, implement, and maintain the core knowledge base: define schemas, normalize entities and units, ensure data quality, and enable efficient retrieval for downstream use.
- Develop and evaluate basic predictive models on top of the extracted data and expose results in a way that is understandable and explainable to users
- Implement simple explainability and provenance: show which passages or records support a given extraction or prediction, and provide clear “why” narratives where possible based on the principles of XAI.
- Prototype ideas end-to-end: from quick experiments and scripts to more robust components that can be integrated into an internal product.
- Collaborate with whoever is working on the interface/backend to expose models and pipelines via clean APIs and iterate based on user feedback.
- Stay up to date with practical developments in LLMs, RAG, and document-understanding, and suggest improvements that can be realistically implemented in a small team.

Requirements

- Strong programming skills in Python, demonstrated through projects, internships, or competitive work (GitHub/portfolio is a plus).
- Hands-on experience with at least one ML/LLM/NLP stack (e.g., PyTorch, TensorFlow, Hugging Face, LangChain, Haystack, or similar).
- Experience working with unstructured text data (parsing PDFs, HTML, or other document formats; building preprocessing/cleaning scripts).
- Basic understanding of core ML concepts: train/validation/test splits, evaluation metrics, overfitting/regularization, feature engineering.
- Familiarity with SQL or NoSQL databases, basic schema design, and writing queries to store and retrieve structured data.
- Comfortable using Linux, command line tools, and Git; able to set up and manage local development environments.
- Mindset: highly self-driven, comfortable multitasking across data, models, and quick experiments; willing to learn whatever is needed to move the project forward in a resource-constrained environment.
- Clear, concise communication in English and the ability to document your work so others can build on it.

Nice-to-have (not mandatory)

- Prior work on RAG, document Q&A, or information extraction into JSON/structured formats.
- Experience running or fine-tuning open-source LLMs locally (e.g., using GPUs, quantized models, or local LLM tools).
- Exposure to basic web APIs (REST) so you can collaborate smoothly with whoever builds the interface.

Qualified and interested candidates may send their curriculum vitae by email to hr@drils.org on or before May 15, 2026 with "**Junior ML Engineer**" as subject line of the email.