Anika Ravindran

NLP Engineer

Boston, MA · anika.ravindran@example.com · (617) 555-0123 · LinkedIn · GitHub

SUMMARY

NLP engineer with 5 years building text classification, information extraction and retrieval systems in Python and PyTorch. Shipped a clause extraction model serving 2 million documents a month and a retrieval-augmented search feature with measured gains in answer accuracy. Comfortable owning the full path from labeling guidelines and evaluation sets to deployment and monitoring.

EXPERIENCE

NLP Engineer — Clausewise Legal Technologies

Boston, MA · Jan 2022 – Present

- Fine-tuned a DeBERTa token classifier for 38 contract clause types, raising macro F1 from 0.71 to 0.86 and serving about 2 million documents a month.
- Built a retrieval-augmented question answering feature over customer contracts using hybrid BM25 and embedding search, lifting answer accuracy on a 600-question test set from 62% to 81%.
- Cut inference cost by 45% by distilling the clause model to a smaller encoder and batching requests on GPU, keeping p95 latency under 300 ms.
- Wrote labeling guidelines and ran agreement checks with 4 legal annotators, reaching a Cohen's kappa of 0.82 before scaling labeling to 15,000 documents.

Machine Learning Engineer, Search — Tradepost Marketplace

Cambridge, MA · Jun 2020 – Dec 2021

- Trained a query classification model on 4 million search logs to route queries to 120 product categories, improving search click-through by 6% in an A/B test.
- Added spelling correction and synonym expansion for listing search, reducing zero-result searches from 9% to 4% of queries.
- Built a nightly evaluation job that tracked precision at 10 across 2,000 labeled queries, catching 3 ranking regressions before release.

SKILLS

NLP: Text classification, NER, information extraction, semantic search, RAG, summarization

Modeling: PyTorch, Hugging Face Transformers, spaCy, scikit-learn, sentence-transformers

Retrieval: Elasticsearch, BM25, FAISS, pgvector, re-ranking

Production: Python, FastAPI, Docker, Kubernetes, AWS SageMaker, MLflow

Evaluation: F1, precision at k, annotation guidelines, inter-annotator agreement, LLM output review

EDUCATION

MS, Computer Science (Natural Language Processing), Northeastern University, 2020

BS, Mathematics, University of Massachusetts Amherst, 2018

Fictional example by FindAJobAI: the person, employers and numbers are invented.