



VIHAAN MEHTA

AI Engineer

Python, retrieval systems, LLM evaluation

Pune, India | vihaan.mehta@example.com | ROLE 02-01

FICTIONAL SAMPLE - all employers, qualifications and outcomes are illustrative and unverified.

EXECUTIVE SUMMARY

AI Engineer with an illustrative 13-year career in analytics, decision support and applied machine learning. Combines python, retrieval systems, llm evaluation with disciplined documentation, practical coordination and clear communication. The sample career progresses from focused execution to independent workstream ownership, with responsibilities and boundaries described for each appointment.

Selected work includes grounded knowledge assistant, document intake assistant and ai evaluation workbench. These examples explain the original problem, individual contribution, deliverables, review approach and remaining limitations rather than relying on unsupported headline claims.

Prepared for experienced ai engineer opportunities requiring dependable delivery, thoughtful professional judgement and collaboration. The qualification narrative includes M.Tech in Artificial Intelligence. Every named organisation and outcome in this record is fictional; professional eligibility is not independently established.

AT A GLANCE

13 years

ILLUSTRATIVE CAREER

4

CAREER APPOINTMENTS

3

WORK EXAMPLES

2

EDUCATION RECORDS

CORE COMPETENCIES

Python / retrieval systems

LLM evaluation / FastAPI

vector search / tool orchestration

prompt testing / data access controls

PROFESSIONAL MISSION

Turn a clear brief into dependable analytics, decision support and applied machine learning work: understand the context, apply python and retrieval systems, record the evidence and explain the limitations before handover.

Career family: AI, Data & Analytics | Snapshot: 2026-09-22

PROFESSIONAL EXPERIENCE / RECENT

AI Engineer

Meridian Analytics Research (fictional)

Jul 2022 - Present | Pune, India

Independent ownership of scoped ai engineer work, coordinating contributors and making review requirements explicit. Includes the first two selected work examples.

- Built retrieval with passage citations and abstention checks.
- Designed reviewed extraction queues and uncertainty flags.
- Led brief clarification and prioritised work using python, retrieval systems and llm evaluation. Raised unresolved constraints before committing to the next stage.
- Coordinated peer reviews and handover preparation; used a data-quality and lineage note to distinguish completed work, assumptions and follow-up needs.
- Supported colleagues with practical examples of fastapi and maintained a concise learning record after important assignments.

Illustrative career responsibilities. Employer confirmation and underlying work records are not supplied.

Senior AI Engineer

Northline Analytics Research (fictional)

Jul 2018 - Jun 2022 | Pune, India

Owned defined assignments and supported cross-functional coordination. Developed deeper practice in llm evaluation and fastapi.

- Created versioned test sets and failure-category reporting.
- Translated incoming requirements into a sequenced plan and aligned responsibilities with the project or service owner.
- Applied vector search and tool orchestration to resolve delivery questions while maintaining source and decision notes.
- Introduced reusable working documents and reviewed exceptions with the responsible specialist rather than silently changing scope.
- Prepared a model or analysis review so the next team could understand the work and remaining questions.

Illustrative career responsibilities. Employer confirmation and underlying work records are not supplied.

PROFESSIONAL WORKING PRINCIPLES

- Define the decision, data lineage and evaluation boundary before choosing a method.
- Separate development, validation and holdout data and record known failure modes.
- Document uncertainty, monitoring ownership and human review for consequential uses.

PROFESSIONAL EXPERIENCE / EARLIER

AI Engineer

Cedarbridge Analytics Research (fictional)

Jul 2015 - Jun 2018 | Pune, India

Progressed from supported tasks to independently managed assignments, with review available for unfamiliar or higher-risk decisions.

- Handled recurring work involving python and retrieval systems using a documented preparation and review process.
- Supported document intake assistant by organising inputs, maintaining issue notes and incorporating reviewer feedback.
- Coordinated colleagues and internal stakeholders using concise status updates, clear questions and agreed next steps.
- Improved record consistency through prompt testing and documented handover expectations.

Illustrative career responsibilities. Employer confirmation and underlying work records are not supplied.

Assistant AI Engineer

Cedarbridge Analytics Research (fictional)

Jul 2013 - Jun 2015 | Pune, India

Built practical foundations through supervised assignments, routine documentation and feedback from experienced colleagues.

- Assisted with python and llm evaluation within an agreed scope and escalated unfamiliar work.
- Prepared inputs and checked completeness before passing work to the responsible reviewer.
- Maintained task records and learned to communicate assumptions, constraints and observed problems clearly.
- Applied review feedback to subsequent assignments and developed a dependable working routine.

Illustrative career responsibilities. Employer confirmation and underlying work records are not supplied.

SELECTED CONTRIBUTIONS

Grounded knowledge assistant. Built retrieval with passage citations and abstention checks. The achievement is the described workstream contribution; independent outcome evidence is not supplied.

Document intake assistant. Designed reviewed extraction queues and uncertainty flags. The achievement is the described workstream contribution; independent outcome evidence is not supplied.

AI evaluation workbench. Created versioned test sets and failure-category reporting. The achievement is the described workstream contribution; independent outcome evidence is not supplied.

SELECTED WORK / 01

Grounded knowledge assistant

P1 | Jan 2023 - Jun 2023 | FICTIONAL WORK EXAMPLE

Context. Fictional internal work programme in analytics, decision support and applied machine learning; not a real client case study.

Problem. Internal answers lacked traceable evidence.

Objective. Create a workable response to this issue through python, explicit review criteria and practical documentation. Agree the boundaries before execution and retain unresolved points for follow-up.

Contribution and ownership. AI Engineer; owned the stated workstream, not the full organisation or every collaborator contribution.

- Built retrieval with passage citations and abstention checks.
- Prepared the scope with the commissioning team, identified unresolved inputs and used python to turn the brief into a sequenced work package.
- Applied retrieval systems and llm evaluation while coordinating reviews with the designated owner. Kept decision notes so collaborators could separate facts, assumptions and changes.
- Assembled the handover material, explained open limitations and agreed which items needed further review rather than presenting them as completed.

Methods. Python / retrieval systems / LLM evaluation / FastAPI

Deliverables. Grounded knowledge assistant - scoped brief; Grounded knowledge assistant - reviewed working package; Grounded knowledge assistant - handover and learning summary

Review. Reviewed the scoped output against the agreed brief, recorded exceptions and checked that key conclusions could be traced to observations. The review package calls for a data-quality and lineage note and a named human reviewer.

Illustrative outcome: the team adopted a repeatable approach for grounded knowledge assistant, with clearer ownership and reviewable records. This is a fictional qualitative result; no real performance measurement or external acceptance evidence is supplied.

Suggested evidence (not supplied): data-quality and lineage note; evaluation protocol; model or analysis review.

SPECIALIST PRACTICE

Evaluation and data stewardship

The sample focuses on python, retrieval systems, llm evaluation. Approvals, supervision and specialist input are identified in each work package; work outside this scope is referred to the designated responsible person.

Peer learning and practical guidance. Created short examples on python and fastapi for colleagues. Separated personal preferences from agreed team procedures and recorded useful questions.

Review and handover discipline. Facilitated practical reviews of document intake assistant, ensuring that unresolved issues retained a named owner rather than disappearing from final presentations.

Examples use synthetic or permissioned data; no claim is made that a model is safe for an untested population.

SELECTED WORK / 02 AND 03

Document intake assistant

P2 | Jan 2024 - Jun 2024 | Meridian Analytics Research (fictional)

Problem. Manual classification delayed incoming requests.

Objective. Create a workable response to this issue through retrieval systems, explicit review criteria and practical documentation. Agree the boundaries before execution and retain unresolved points for follow-up.

Ownership. AI Engineer; owned the stated workstream, not the full organisation or every collaborator contribution.

- Designed reviewed extraction queues and uncertainty flags.
- Prepared the scope with the commissioning team, identified unresolved inputs and used llm evaluation to turn the brief into a sequenced work package.
- Applied fastapi and vector search while coordinating reviews with the designated owner. Kept decision notes so collaborators could separate facts, assumptions and changes.
- Assembled the handover material, explained open limitations and agreed which items needed further review rather than presenting them as completed.

Methods. retrieval systems / LLM evaluation / FastAPI / vector search

Delivered. Document intake assistant - scoped brief; Document intake assistant - reviewed working package; Document intake assistant - handover and learning summary

Illustrative outcome: the team adopted a repeatable approach for document intake assistant, with clearer ownership and reviewable records. This is a fictional qualitative result; no real performance measurement or external acceptance evidence is supplied.

Suggested evidence (not supplied): data-quality and lineage note; evaluation protocol; model or analysis review.

AI evaluation workbench

P3 | Jan 2019 - Jun 2019 | Northline Analytics Research (fictional)

Problem. Model changes were judged through anecdotal examples.

Objective. Create a workable response to this issue through llm evaluation, explicit review criteria and practical documentation. Agree the boundaries before execution and retain unresolved points for follow-up.

Ownership. AI Engineer; owned the stated workstream, not the full organisation or every collaborator contribution.

- Created versioned test sets and failure-category reporting.
- Prepared the scope with the commissioning team, identified unresolved inputs and used vector search to turn the brief into a sequenced work package.
- Applied tool orchestration and prompt testing while coordinating reviews with the designated owner. Kept decision notes so collaborators could separate facts, assumptions and changes.
- Assembled the handover material, explained open limitations and agreed which items needed further review rather than presenting them as completed.

Methods. LLM evaluation / FastAPI / vector search / tool orchestration

Delivered. AI evaluation workbench - scoped brief; AI evaluation workbench - reviewed working package; AI evaluation workbench - handover and learning summary

Illustrative outcome: the team adopted a repeatable approach for ai evaluation workbench, with clearer ownership and reviewable records. This is a fictional qualitative result; no real performance measurement or external acceptance evidence is supplied.

Suggested evidence (not supplied): data-quality and lineage note; evaluation protocol; model or analysis review.

EDUCATION AND SPECIALIST LEARNING

M.Tech in Artificial Intelligence

Asterbridge Institute of Professional Studies (fictional institution)
Jul 2011 - May 2013 | completed - fictional record
Focus: Python, retrieval systems, FastAPI. Specialist study on grounded knowledge assistant; an illustrative learning project, not a published result.

B.Tech in Computer Science

Cedarhaven College of Applied Studies (fictional institution)
Jul 2007 - May 2011 | completed - fictional record
Focus: LLM evaluation, vector search, tool orchestration. Applied coursework in python, documentation and reviewed practical assignments.

CONTINUING PROFESSIONAL DEVELOPMENT

Experiment design practicum (2023). Python and fastapi in the context of ai engineer work. Used reflective exercises and a bounded practice example related to grounded knowledge assistant.

Responsible data handling workshop (2024). Retrieval systems and vector search in the context of ai engineer work. Used reflective exercises and a bounded practice example related to document intake assistant.

Model monitoring case lab (2025). Llm evaluation and tool orchestration in the context of ai engineer work. Used reflective exercises and a bounded practice example related to ai evaluation workbench.

Provider for these illustrative learning records: Meridian Professional Learning Studio (fictional). Not a licence or certification.

LANGUAGES AND WORKING PREFERENCES

English and Hindi - professional working proficiency in this fictional record.

Work arrangement: hybrid; remote feasibility discussed per role. Relocation: open to discussion - sample preference. Availability: not confirmed; discuss before an interview.

CREDENTIAL AND EVIDENCE STATUS

Examples use synthetic or permissioned data; no claim is made that a model is safe for an untested population.

Synthetic qualification and institution names illustrate profile fields only; they are not a validated qualification route or recognised accreditation claim.

No real referee, endorsement, award, publication, membership or licence is supplied.

Use for templates, product demonstrations and test data only. Replace fictional claims and remove the demo label only after approval of real candidate information.

Contact placeholder: vihaan.mehta@example.com. No phone, live profile URL or verified reference is supplied.

Complete machine-readable data: profile.json. Interactive companion: portfolio.html. Reference CV used for structure only; this is not Ajay Prajapat's biography.