

Skills

Python

Applied to grounded knowledge assistant through documented preparation, execution and review.

retrieval systems

Applied to document intake assistant through documented preparation, execution and review.

LLM evaluation

Applied to ai evaluation workbench through documented preparation, execution and review.

FastAPI

Applied to grounded knowledge assistant through documented preparation, execution and review.

vector search

Applied to document intake assistant through documented preparation, execution and review.

tool orchestration

Applied to ai evaluation workbench through documented preparation, execution and review.

prompt testing

Applied to grounded knowledge assistant through documented preparation, execution and review.

data access controls

Applied to document intake assistant through documented preparation, execution and review.

Certifications

Experiment design practicum

Meridian Professional Learning Studio (fictional) 2023 continuing learning - not a licence or certification 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

Meridian Professional Learning Studio (fictional) 2024 continuing learning - not a licence or

Vihaan Mehta

AI Engineer · FICTIONAL SAMPLE

Pune, India

Summary

Fictional demonstration candidate. All employers, education, projects, achievements and outcomes are illustrative and unverified. Not a real application or licence record. 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. 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.

Experience

AI Engineer - Meridian Analytics Research (fictional)

2022-07 - Present

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)

2018-07 - 2022-06

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.

AI Engineer - Cedarbridge Analytics Research (fictional)

2015-07 - 2018-06

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.

certification 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

Meridian Professional Learning Studio (fictional) 2025 continuing learning - not a licence or certification 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.

Languages

English

professional working - fictional sample

Hindi

professional working - fictional sample

Vihaan Mehta

AI Engineer · FICTIONAL SAMPLE

Pune, India

Assistant AI Engineer - Cedarbridge Analytics Research (fictional)

2013-07 - 2015-06

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.

Education

M.Tech in Artificial Intelligence - Asterbridge Institute of Professional Studies (fictional institution)

2011-07 - 2013-05

completed - fictional record 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)

2007-07 - 2011-05

completed - fictional record LLM evaluation vector search tool orchestration Applied coursework in python, documentation and reviewed practical assignments.

Projects

Grounded knowledge assistant

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. 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. Limitations: Synthetic demonstration only. Examples use synthetic or permissioned data; no claim is made that a model is safe for an untested population. No underlying client documents, independently verified measurements or signed approval records are attached.

Document intake assistant

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

Vihaan Mehta

AI Engineer · FICTIONAL SAMPLE

Pune, India

before execution and retain unresolved points for follow-up. 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 Deliverables: Document intake assistant - scoped brief; Document intake assistant - reviewed working package; Document intake 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 evaluation protocol and a named human reviewer. 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. Limitations: Synthetic demonstration only. Examples use synthetic or permissioned data; no claim is made that a model is safe for an untested population. No underlying client documents, independently verified measurements or signed approval records are attached.

AI evaluation workbench

Fictional internal work programme in analytics, decision support and applied machine learning; not a real client case study. 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. 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 Deliverables: AI evaluation workbench - scoped brief; AI evaluation workbench - reviewed working package; AI evaluation workbench - 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 model or analysis review and a named human reviewer. 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. Limitations: Synthetic demonstration only. Examples use synthetic or permissioned data; no claim is made that a model is safe for an untested population. No underlying client documents, independently verified measurements or signed approval records are attached.

Achievements

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



Vihaan Mehta

AI Engineer · FICTIONAL SAMPLE
Pune, India

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

