N/39/2026-NeGD
National e-Governance Division
Digital India Corporation
Electronics Niketan, 6, CGO Complex Lodhi Road,
New Delhi – 110003
Website: www.negd.gov.in / www.dic.gov.in

Web Advertisement
22.10.2026

The National e-Governance Division (NeGD) is an independent business division under the Digital India Corporation, Ministry of Electronics and Information Technology. NeGD has been playing a pivotal role in supporting MeitY in Programme Management and implementation of e-Governance projects and initiatives undertaken by various Ministries/ Departments, both at the Central and State levels.

NeGD has been spearheading several innovative initiatives under the aegis of the Digital India Programme. Those have been developed keeping the vision areas of Digital India at the core- providing digital infrastructure as a core utility to every citizen, governance and services on demand and in particular, digital empowerment of the citizens of our country; some of these initiatives include DigiLocker, UMANG, Poshan Tracker, OpenForge Platform, API Setu, National Academic Depository, Academic Bank of Credits, Learning Management System.

It has myriad roles and responsibilities from supporting Central Line Ministries and State Government Departments on e-Governance projects, reviewing State Action Plans, offering support in technology management, strategy formulation & implementation of Emerging Technologies viz. AI, Blockchain, GIS etc., to facilitating digital diplomacy with focus on Indian startups and products

NeGD has been a leader in implementation and execution of a gamut of pilot/ infrastructure/ technical/ special projects and support components to framing core policies, project appraisals, R&D, and guiding /conducting assessments, undertaking activities for building capacities of both Government officials and] other stakeholders, and creating mass awareness about schemes and services under the Digital India Programme.

NeGD is currently inviting applications for the following positions purely on Contract basis initially for a period of 3 years which is further extendable as per the requirement of the project.

S. No	Position	Vacancy
1	Principal Architect - AI	01
2	MLOPs Lead	01
3	MLOPs Engineer	01
4	AI Safety Researcher	01
5	AI/ML Lead	04
6	Data Science Lead	02

* The maximum age limit shall be 55 years on the closing date of receipt of application.
** The place of posting shall be in New Delhi but transferable to project locations of NeGD as per existing policy of NeGD/DIC.
Screening of applications will be based on qualifications, age, and relevant experience. NeGD reserves the right to fix higher threshold of qualifications and experience for screening and limiting the number of candidates for interview. Only shortlisted candidates shall be invited for interviews. NeGD reserves the right to not to select any of the candidates without assigning any reason thereof.
The details can be downloaded from the official website of **DIC** and **NeGD**, viz. **www.dic.gov.in**, **www.negd.gov.in**.

Eligible candidates may apply ONLINE: https://ora.digitalindiacorporation.in/

Last date for submission of applications will be: 13.11.2026

Designation: Principal Architect - AI

Educational Qualification

- B.Tech / M.Tech / MS / Ph.D. in Computer Science, Information Technology, Artificial Intelligence, or a related technical field.
- Certifications preferred: AWS Solutions Architect Professional, Azure Solutions Architect Expert, or GCP Professional Cloud Architect.
- Research papers, case studies, or significant open-source contributions (are preferred)

Experience

- Minimum 12+ years of total professional experience.
- At least 5 years in AI/ML system design
- Proven experience designing and deploying multi-component AI solutions for government or enterprise environments.

Key Responsibilities

- Design and implement **end-to-end solution architectures** embedding AI/ML capabilities into NeGD and allied government applications.
- Define integration standards, APIs, and interoperability frameworks between AI systems and existing Digital India platforms.
- Oversee data flow architecture, security, privacy, and Responsible AI implementation across solutions.
- Architect and oversee integration of AI applications with NeGD's current enterprise applications and digital platforms.
- Evaluate and recommend cloud-native and on-premises deployment models in alignment with MeitY guidelines.
- Provide architectural leadership, ensuring scalability, maintainability, and compliance with MeitY and NIC technical standards.
- Collaborate with Data Science, MLOps, and Engineering teams for unified architecture governance.
- Conduct periodic architectural reviews and provide technical mentoring to development teams.

Technical Competencies:

- **Architecture:** Microservices and serverless patterns, event-driven and domain-driven design, high-level data & model lifecycle architecture
- **Cloud:** AWS (Bedrock, SageMaker, Lambda, API Gateway, S3), Azure (OpenAI Service, ML Studio, AKS), GCP (Vertex AI, Big Query, Cloud Storage), multi-cloud integration & cost optimization
- **AI Frameworks:** TensorFlow, PyTorch, Hugging Face Transformers, Lang Chain (LLM orchestration), prompt-engineering and responsible AI practices
- **Data & Storage:** SQL (PostgreSQL, MySQL), NoSQL (MongoDB, DynamoDB), vector databases (Pinecone, Weaviate) for RAG & semantic search
- **APIs & Integration:** REST, GraphQL, WebSocket for real-time chat/voice
- **Conversational AI & Voice:** Speech-to-Text and Text-to-Speech APIs (AWS Transcribe/Polly, Azure Speech), RAG/chatbot pipeline design with vector stores & LLMs
- **Security & Governance:** IAM and RBAC across clouds, encryption in transit & at rest, data privacy & AI governance (GDPR, SOC2)

Designation: MLOPs Lead

Educational Qualification

- B.Tech / M.Tech / M.S. in Computer Science, Data Engineering, AI or related discipline.
- Certification in cloud DevOps or MLOps platforms (AWS DevOps Engineer, Azure DevOps Expert, GCP Professional ML Engineer) is highly desirable.
- Contributions to MLOps or DevOps open source projects is preferred.

Experience:

- 7–10 years in machine learning operations or DevOps engineering.
- Minimum 4 years building CI/CD pipelines for AI/ML model deployment in enterprise or government ecosystems.
- Proven experience with containerized and microservice architectures.

Key Responsibilities:

- Design and manage continuous integration and delivery (CI/CD) pipelines for AI/ML models across multiple environments.
- Establish model versioning, deployment, monitoring, and rollback mechanisms to ensure stability and traceability.
- Automate training, testing, and serving workflows using containerized solutions.
- Define infrastructure-as-code templates for scalable AI deployment on on-prem or cloud environments.
- Collaborate with Data Science and Engineering teams to standardize model input/output formats and performance metrics.
- Implement logging, monitoring, and alerting for deployed models to ensure high availability and accuracy over time.
- Ensure compliance with Responsible AI guidelines for deployment, including bias auditing and explainability tracking.

Technical Competencies:

- MLOps Platforms: MLflow, Kubeflow, Azure ML, AWS SageMaker Pipelines, GCP Vertex AI Pipelines for end-to-end ML workflow orchestration
- Containerization: Docker, Kubernetes, Helm charts, container registries, and microservices architecture for ML workloads
- CI/CD: Jenkins, GitLab CI, GitHub Actions, Azure DevOps with specialized ML pipeline integration and automated testing
- Infrastructure-as-Code: Terraform, CloudFormation, Ansible for reproducible ML infrastructure provisioning and management
- Cloud Platforms: AWS (EKS, Lambda, ECR, S3), Azure (AKS, Container Registry, Blob Storage), GCP (GKE, Cloud Build, Cloud Storage)
- Model Serving: TorchServe, TensorFlow Serving, Seldon, KServe, REST APIs, and real-time inference infrastructure.
- Programming Languages: Python for automation, Bash scripting, YAML for configuration management, basic understanding of Go/Java
- Database & Storage: Feature stores (Feast, Tecton), model registries, data versioning (DVC), and distributed storage systems
- Workflow Orchestration: Apache Airflow, Prefect, Argo Workflows for complex ML pipeline scheduling and dependency management

Designation: MLOPs Engineer

Educational Qualification

- B. Tech / M.Tech in Computer Science, AI/ML, or related discipline.
- Certifications in DevOps or Cloud Infrastructure (AWS, Azure, or GCP) preferred.
- Research papers, case studies, or significant open-source contributions (are preferred)

Experience:

- 3–6 years in operationalizing AI/ML models with proven experience in CI/CD automation.
- Prior exposure to deploying ML pipelines for NLP, computer vision, or speech systems.
- Familiarity with monitoring, model lifecycle management, and performance logging.

Key Responsibilities:

- Deploy and manage AI/ML models in development, staging, and production environments.
- Build and maintain automated pipelines for continuous integration and delivery (CI/CD).
- Implement real-time monitoring for model drift, latency, and inference performance.
- Collaborate with Solution Architect and MLOps Lead to standardize deployment infrastructure.
- Ensure reproducibility, rollback, and version control for deployed models.
- Integrate AI services with NeGD's standard APIs and observability frameworks.
- Maintain deployment logs, error reports, and environment snapshots for audit readiness.

Technical Competencies:

- Infrastructure Tools: Jenkins, GitLab CI/CD, Docker, Kubernetes.
- Monitoring & Logging: Prometheus, Grafana, ELK Stack.
- ML Lifecycle Management: MLflow, Kubeflow, DVC.
- Cloud Platforms: AWS SageMaker, Azure ML Studio, GCP Vertex AI.
- Core: Python, Bash scripting, YAML/JSON configuration, Linux systems.
- CI/CD: Jenkins, GitLab CI, or GitHub Actions for automated deployments,Terraform
- Governance: Traceability and Responsible AI compliance in deployment

Designation: AI Safety Researcher

Educational Qualification

- B.Tech./B.E. or M.Tech./M.S./M.Sc. in Computer Science, Information Security, AI/ML, or a related quantitative discipline **(Must have)**
- Advanced degree (M.Tech./M.S./Ph.D.) with a thesis or published work in adversarial ML, AI security, LLM safety, or Responsible AI is highly desirable
- Certifications (Desirable): OSCP, GWAPT or CEH combined with demonstrable AI/ML security work; DeepLearning.AI or equivalent ML foundations; MLSecOps or LLM-security specialist certifications where available
- Non-traditional backgrounds with demonstrable adversarial ML research, published safety work, credible LLM red-team disclosures, or CTF/red-team achievements will be considered in lieu of formal qualification

Experience

- 6+ years total in ML, applied AI, security research, or a closely related discipline; **minimum 3 years specifically in adversarial ML, AI red teaming, LLM safety evaluation, or AI/ML security research**
- Demonstrable hands-on LLM red-teaming experience with documented prompt-injection, jailbreak, indirect-prompt-injection or agentic-tool-misuse campaigns against production or production-like systems
- Demonstrable adversarial ML work — evasion, model inversion, membership inference, model extraction, or data poisoning — against non-toy classifiers, vision models, or NLP systems
- Prior experience delivering safety, red-team, or Responsible AI work in government, BFSI, healthcare, or another regulated sector is a strong plus
- Prior experience advising or upskilling non-specialist engineering, security, or compliance teams on AI-specific security concerns is desirable

Key Responsibilities

Adversarial ML & Model Security Testing

- Design and execute adversarial attack campaigns — evasion, model inversion, membership inference, model extraction, data poisoning — against pod deliverables including document intelligence, predictive-analytics and fraud-scoring models, and visual AI, face-embedding and liveness models
- Develop and maintain a reusable adversarial ML attack library and test harness for repeatable pre-production evaluation
- Produce adversarial-robustness reports for each pod service with quantitative measures, reproducible attack notebooks, and prioritised mitigation guidance

LLM, RAG & Agentic AI Red Teaming

- Lead structured red-team exercises against LLM, RAG and agentic AI deployments across the programme.
- Cover prompt injection, jailbreaks, indirect prompt injection via retrieved documents, data exfiltration, unsafe tool invocation, sandbox escape and policy-boundary violations by agents
- Develop and maintain red-team playbooks tailored to government AI use cases; publish anonymised playbooks and evaluation sets to AIKosh under standard metadata

- Advise pods on guardrail selection, output filtering, RAG source-integrity controls, retrieval provenance, and agent policy design

Hallucination, Calibration & Responsible AI Evaluation

- Design and run hallucination measurement, groundedness checks, and calibration/uncertainty evaluation for RAG systems and generative outputs across teams
- Run bias and fairness audits using quantitative metrics — demographic parity, equalised odds, calibration, subgroup accuracy.
- Conduct explainability evaluation (SHAP, LIME, Captum) and lightweight privacy impact assessments on team deliverables
- Author the technical safety evaluation content in model cards, dataset sheets, bias/hallucination/safety evaluation reports, and Responsible AI evidence per the MeitY Responsible AI advisory and the IndiaAI Safe & Trusted AI framework.

AI Security Architecture & Threat Modelling

- Own AI-specific threat modelling end-to-end for all pod systems under development — STRIDE, MITRE ATT&CK, MITRE ATLAS and OWASP Top 10 for LLMs — including data pipelines, retrieval sources, model artefacts, prompt paths, tool interfaces and output surfaces
- Review deployment configurations across IndiaAI Compute, on-premise/sovereign GPU and CSP-API paths for AI-specific security risks; recommend controls and defensive patterns
- Contribute AI security and Responsible AI requirements to NeGD RFQs, RDRs and procurement documents; review empanelled-agency proposals for AI security posture

Cross-Programme Advisory, Upskilling & Governance Input

- Advise each team's AI QA Engineer on safety and Responsible AI test design, sample selection, evaluation metrics and evidence capture
- Brief and upskill teams on AI-specific security concerns and Responsible AI controls.

Research, Publication & Knowledge Transfer

- Track adversarial ML, LLM safety, agentic-AI safety and Responsible AI research literature; translate relevant findings into team-usable checks, controls and evaluation additions
- Maintain the programme's open-source AI safety testing toolkit; deposit reusable notebooks, evaluation harnesses and playbooks to OpenForge, and publish reusable evaluation sets and safety artefacts to AIKosh under standard metadata for reuse

Technical Competencies

- **Programming & ML Frameworks:** Python (advanced); PyTorch or TensorFlow; Hugging Face Transformers; standard data-science tooling (NumPy, pandas, scikit-learn)
- **Adversarial ML:** Adversarial Robustness Toolbox (ART), Foolbox, CleverHans or equivalent; ability to implement custom attacks and defences; knowledge of certified-robustness techniques
- **LLM & Agentic AI Red Teaming:** Demonstrable production-relevant experience with prompt injection, jailbreak, indirect-prompt-injection, data-exfiltration, tool-misuse, sandbox-escape and multi-turn manipulation; familiarity with LLM guardrail frameworks (NeMo Guardrails, Guardrails AI, Llama Guard) and open red-team datasets
- **Hallucination, Calibration & RAG Evaluation:** RAGAS or equivalent;

groundedness metrics, faithfulness scoring, retrieval quality metrics, calibration and uncertainty quantification; ability to build custom evaluation harnesses for RAG and generative pipelines

- **Threat Modelling & AI Security Frameworks:** STRIDE, MITRE ATT&CK, MITRE ATLAS, OWASP Top 10 for LLMs, OWASP ML Top 10; ability to translate threat models into control specifications and test cases
- **Explainability & Fairness:** SHAP, LIME, Captum for model explainability; Fairlearn and AI Fairness 360 for fairness metrics; ability to design subgroup-fairness protocols for identity verification, fraud and predictive models
- **Privacy-Enhancing Techniques:** Working awareness of differential privacy, federated learning, PII redaction and anonymisation techniques; ability to run privacy impact assessments on model and data pipelines
- **Government Frameworks & Standards:** IndiaAI Safe & Trusted AI framework, MeitY Responsible AI advisory, MeitY Security Policy and Guidelines, CERT-In directions, DPDPA 2023
- **Communication & Advisory:** Ability to author clear technical safety reports for a mixed engineering, architecture and executive audience; ability to brief and upskill non-AI-security engineers and compliance colleagues; ability to represent the programme in cross-ministerial working groups

Designation: AI/ML Lead

Specialisation	Positions	Focus
NLP & Document AI	1	OCR and document parsing, information extraction, PII redaction, and retrieval-augmented question answering
Speech AI	1	Speech recognition and synthesis for Indian languages, and voice and chat assistants
Computer Vision	1	Face verification and liveness, document forgery detection, and image and geospatial analytics
Agentic AI	1	LLM agents that complete multi-step work in enterprise systems under human approval

Role Summary

The AI/ML Lead will lead an AI team in one specialisation, owning the design, development, evaluation and deployment of reusable AI services and providing technical oversight of vendor teams. The role requires hands-on depth in the specialisation, rigorous evaluation, and the judgement to build services once for reuse across government departments, using open-weight and Indian models hosted on IndiaAI compute.

Educational Qualification

- B.Tech./B.E., M.Tech./M.S. or M.Sc. in Computer Science, AI/ML, Data Science, Electronics or a related discipline
- Advanced degree (M.Tech./M.S./Ph.D.) or published work in the chosen specialisation is highly desirable
- Certifications are not mandatory; demonstrated work, as listed under Minimum Evidence at Screening, shall carry greater weight
- Candidates from non-traditional backgrounds with strong verifiable work will be considered in lieu of formal qualification

Experience

- 7+ years developing and deploying AI/ML systems (10+ years preferred); **minimum 4 years building systems in the chosen specialisation deployed in production** (Agentic AI: minimum 2 years building LLM applications or agents in production)
- At least 2 years leading a team of engineers or data scientists, including setting technical direction and reviewing their work
- **Demonstrable delivery of at least one system in the chosen specialisation used by real users at scale**, with documented evaluation: field-level accuracy (NLP & Document AI), word error rate or resolution rate (Speech AI), error rates at operating thresholds (Computer Vision), or task-success rates and safety controls (Agentic AI)
- Hands-on experience fine-tuning open-weight or foundation models on domain data
- Experience with Indian languages, scripts or code-mixed data is a strong plus
- Experience overseeing vendor delivery is desirable

Key Responsibilities

Technical Leadership

- Own the technical roadmap of the team's reusable services, choosing between prompting, retrieval, fine-tuning and smaller specialised models on accuracy, cost, latency and data-governance grounds, and documenting the rationale

- Lead in-house engineers and data scientists, review their work, and remain hands-on in code, training and debugging

Specialisation Responsibilities

- **NLP & Document AI:** OCR, layout analysis and parsing for scanned, handwritten and mixed-script documents; extraction, classification, PII redaction, summarisation, similarity and de-duplication; retrieval-augmented question answering over rules, circulars and case files, with citations and grounding checks
- **Speech AI:** speech recognition, synthesis, translation and language identification for Indian languages, including 8 kHz telephone audio and code-mixed speech; assistants on web, WhatsApp, IVR and mobile with retrieval-grounded answers, human hand-off and safe failure; real-time voice pipelines within latency budgets
- **Computer Vision:** face detection, 1:1 verification, 1:N search, liveness and presentation-attack detection at large scale; tamper, forgery and photo-substitution detection for identity documents and certificates; field-photo and satellite imagery analytics; deployment on server, mobile and edge
- **Agentic AI:** agents for planning, tool use, memory and multi-agent coordination using LangGraph or equivalent and the Model Context Protocol (MCP); least-privilege tool permissions, human approval for every consequential action, audit trails and rollback; reusable agent templates for drafting, document scrutiny and request routing

Evaluation & Safety

- With the QA Lead, build golden datasets and specialisation-appropriate metrics, set release thresholds, and run subgroup analysis across languages, document types, demographics or acoustic conditions
- With the AI Safety team, test services for prompt injection, presentation attacks, adversarial inputs and leakage of personal data before release

Deployment & Reuse

- Deploy services through the AI platform and gateway, publish them to the service catalogue with documentation and SLAs, and contribute models and evaluation sets to AIKosh

Vendor Oversight

- Review designs, code and evaluation evidence of vendor teams, and recommend acceptance or rejection of deliverables

Responsible AI & Privacy

- Apply anonymisation, data minimisation and DPDPA 2023 requirements, with stricter handling for biometric data in Computer Vision, and document models through model cards and evaluation records

Technical Competencies

- **Common:** Python (expert level), PyTorch, Hugging Face Transformers, PEFT and TRL (LoRA, QLoRA, DPO, GRPO); vLLM or SGLang; MLflow, Docker, Git; evaluation design with bootstrap confidence intervals; Langfuse or Arize Phoenix
- **Indic Language AI:** AI4Bharat models (IndicConformer, IndicTrans2, IndicBERT, Indic Parler-TTS), Sarvam models and Bhashini services; code-mixed and transliterated text
- **NLP & Document AI:** document-parsing models such as PaddleOCR-VL, MinerU or GLM-OCR, and PP-OCR; Docling; vision-language models such as the Qwen-VL family; hybrid search and re-ranking; pgvector, Qdrant or Milvus; RAGAS or DeepEval; Microsoft Presidio
- **Speech AI:** IndicConformer, Whisper, NVIDIA Parakeet or Canary; NVIDIA NeMo or ESPnet; Silero VAD, pyannote; WebRTC, SIP, Pipecat or LiveKit Agents; WhatsApp Business Platform; WER and CER tooling

- **Computer Vision:** YOLO26 or YOLO11, RF-DETR, SAM 3, DINOv3, SigLIP 2; ArcFace or AdaFace with InsightFace; ISO/IEC 30107-3 presentation-attack metrics; image-manipulation detection (such as TruFor); ONNX Runtime, TensorRT, LiteRT
- **Agentic AI:** LangGraph, LlamaIndex, OpenAI Agents SDK or Google ADK; MCP; Temporal; Open Policy Agent; trajectory and tool-use evaluation (Inspect, DeepEval or LangSmith); prompt-injection testing for agents (such as AgentDojo)
- **Governance:** DPDPA 2023, MeitY Responsible AI advisory, IndiaAI Safe & Trusted AI framework, OWASP Top 10 for LLM Applications

Designation: Data Science Lead

Specialisation	Positions	Focus
Predictive Analytics	1	Forecasting, risk and propensity scoring, and early-warning models
Fraud & Risk Analytics	1	Fraud and anomaly detection, entity resolution and network analysis on claims and transactions

Role Summary

The Data Science Lead will lead a data science team in one specialisation, owning the predictive, risk and anomaly models that government departments use for operational decisions, and providing technical oversight of vendor teams. The role requires statistical depth, sound judgement on the choice of method, rigorous validation with quantified uncertainty, and the ability to explain models to senior decision makers and investigators.

Educational Qualification

- B.Tech./B.E., M.Tech./M.S. or M.Sc. in Statistics, Mathematics, Economics, Computer Science, Data Science or a related quantitative discipline
- Ph.D. or published applied work in statistics, econometrics, forecasting or fraud analytics is highly desirable
- Certifications are not mandatory; demonstrated work, as listed under Minimum Evidence at Screening, shall carry greater weight
- Candidates with strong verifiable work (Kaggle Master level or above, published research or significant open-source contributions) will be considered in lieu of formal qualification

Experience

- 8+ years in data science, statistics or applied machine learning (12+ years preferred); **minimum 4 years building models in the chosen specialisation used in operational decisions**
- At least 3 years leading data science teams or projects, including review of other data scientists' models
- **Demonstrable ownership of at least two predictive, risk or fraud systems in production**, with documented validation and measured outcomes
- Predictive Analytics: experience with hierarchical, spatial or panel data, and with missing, delayed or inconsistent data
- Fraud & Risk Analytics: experience detecting fraud or anomalies with few labels, including record linkage across sources and graph-based detection
- Experience with causal inference or programme evaluation is a strong plus; experience in the public sector, health, agriculture, banking, insurance or payments is a plus

Key Responsibilities

Problem Formulation

- Translate business questions into well-defined problems with a clear decision, baseline and evaluation metric
- Assess data feasibility, including coverage, reporting delays, label quality, leakage and bias, and advise where the data is insufficient for the objective

Specialisation Responsibilities

- **Predictive Analytics:** time-series, hierarchical and spatial forecasts at district, state and national level, with reconciliation across levels and backtesting; risk, propensity and prioritisation scores; early-warning thresholds designed with the operational teams who will use them
- **Fraud & Risk Analytics:** fraud typologies defined with investigators; rules combined with unsupervised, semi-supervised and positive-unlabelled learning; entity resolution across individuals, bank accounts, addresses and documents; graph analysis of linked cases; risk scores with reason codes and review queues sized to investigator capacity; checks against identity, tax, document and bank-account records

Validation, Uncertainty & Monitoring

- Validate models with leakage-free schemes, and quantify uncertainty through prediction intervals, calibration and conformal prediction
- Monitor drift after go-live, including adaptation by fraudsters, and assess whether deployed models improved outcomes using experiments or quasi-experimental methods

Responsible AI & Governance

- Audit fairness across regions and population groups, provide explanations for decisions affecting individuals, and maintain model cards, evaluation records and audit trails
- Apply anonymisation and data minimisation in compliance with DPDPA 2023

Analytics Products

- Deliver dashboards and analytical services on a shared metric layer, with the Data Engineering Lead and the engineering team

Team Leadership & Vendor Oversight

- Lead in-house data scientists, review their work, and remain hands-on
- Provide technical oversight of vendor teams, and recommend acceptance or rejection of deliverables

Technical Competencies

- **Statistics & Causal Inference:** generalised linear, hierarchical and Bayesian models (PyMC); A/B testing, difference-in-differences, synthetic control, regression discontinuity (DoWhy, EconML)
- **Machine Learning:** XGBoost, LightGBM and CatBoost; tabular foundation models (TabPFN); explainability (SHAP); calibration
- **Uncertainty:** prediction intervals, conformal prediction (MAPIE), bootstrap methods
- **Predictive Analytics:** ARIMA, ETS and state-space models; hierarchical reconciliation (MinT); statsforecast, Darts or GluonTS; time-series foundation models (Chronos-2, TimesFM); GeoPandas and H3
- **Fraud & Risk Analytics:** PyOD, Isolation Forest, autoencoders; Splink, Zingg or Dedupe with transliteration-aware matching of Indian names and addresses; Neo4j, NetworkX, PyTorch Geometric or DGL; weak supervision and Cleanlab; precision at k and cost-sensitive metrics
- **Data & Engineering:** advanced SQL, PySpark, Polars or DuckDB; Python with pytest, Git, MLflow and DVC
- **Communication:** Superset, Power BI or Tableau; explaining limitations and uncertainty to non-technical audiences
- **Governance:** Fairlearn; DPDPA 2023; MeitY Responsible AI advisory

General Conditions applicable to all applicants covered under this advertisement

1. Those candidates, who are already in regular or contractual employment under Central / State Government, Public Sector Undertakings or Autonomous Bodies, are expected to apply through proper channel or attach a 'No Objection Certificate' from the employer concerned with the application OR produce No Objection Certificate at the time of interview.
2. NeGD reserves the right to fill all or some or none of the positions advertised without assigning any reason as it deems fit.
3. The positions are purely temporary in nature for the project of NeGD/DIC and the appointees shall not derive any right or claim for permanent appointment at NeGD/DIC or on any vacancies existing or that shall be advertised for recruitment by NeGD in future.
4. NeGD reserves the right to terminate the appointments of all positions with a notice of one month or without any notice by paying one month's salary in lieu of the notice period.
5. The designation of the selected candidates shall be mapped as per the existing designation policy of NeGD.
6. In case of a query, the following officer may be contacted:

HR Team
National e Governance Division, 4th Floor, Electronics Niketan,
6-CGO, Complex Lodhi Road, New Delhi – 110003
Email: Negdhr@digitalindia.gov.in