Learn how to move from research notebooks to reliable, monitored machine learning services in bioinformatics. This module covers MLOps foundations, CI and CD, model registry, packaging and deployment patterns plus monitoring and drift management for omics and clinical ML pipelines.

MLOps for Bioinformatics — CI/CD, Model Registry & Monitoring

Help Desk · WhatsApp

Session Index

Session 1 — MLOps Foundations & Project Structuring	Session 2 — CI/CD for Data & Models
Session 3 — Model Registry, Packaging & Deployment	Session 4 — Monitoring, Drift & Operations Governance

Session 1
Fee: Rs 8800 Apply Now

MLOps Foundations & Project Structuring

From research notebooks to MLOps mindset

pain points in hand run pipelines | separation of research vs production concerns | basic MLOps lifecycle for bioinformatics

Structuring ML projects for teams

standard folders for data, code and configs | config

files and run parameters simple conventions for scripts and notebooks

Version control and issue tracking in practice

Git branching strategies for ML repos code review and pull requests simple backlog and issue templates

Session 2
Fee: Rs 11800 Apply Now

CI/CD for Data & Models

Continuous integration for ML repos

lightweight tests for data and code linting and style checks automated reports on pull requests

Continuous delivery patterns for models

build pipelines for training jobs test and staging environments promotion rules to production

Automating data and feature checks

schema and range checks for inputs simple distribution shift alerts at build time capturing build artefacts for traceability

Session 3
Fee: Rs 14800 Apply Now

Model Registry, Packaging & Deployment

Using a model registry effectively

registering models with metadata tracking versions, metrics and lineage staging, production and archived states

Packaging models as services or jobs

container images for inference batch scoring vs online scoring simple API or CLI patterns

Deployment considerations in bioinformatics settings

on premises vs cloud constraints resource limits for heavy models simple blue green or canary style rollouts

Session 4
Fee: Rs 18800 Apply Now

Monitoring, Drift & Operations Governance

Setting up monitoring for ML in production

logging predictions and inputs tracking latency and resource usage basic alerts and dashboards

Data drift, concept drift and model decay

simple drift detection statistics periodic back testing of performance triggers for retraining and rollback

Governance, approvals and runbooks

defining roles and responsibilities change management and approval steps operations runbooks and incident logs