iiit-b International Institute of Information Technology Bangalore
ज्ञानमुत्तमम्

powered by upGrad

Executive Diploma in

Machine Learning & Artificial Intelligence

With MLOps and
Generative AI Specialisations

Duration 12-15 months

Largest ML & AI program in India

with 9+ Years of Legacy &
10,000+ learners

powered by upGrad

Glimpses From Convocation Ceremony

About IIITB

The International Institute of Information Technology Bangalore (IIITB)

Established in 1998, IIITB is a premier institute known for its interdisciplinary approach, integrating technology with social sciences. Supported by the Government of Karnataka and the IT industry, IIITB fosters strong academic-industry partnerships and attracts top talent from across India and abroad through its merit-based selection process.

The institute has graduated over 3,500 students, many of whom work at leading IT companies globally. With a focus on research and development in fields like Artificial Intelligence (AI) and Machine Learning (ML), IIITB is recognised as a leader in AI education.

Ranked 74th in the Engineering category of the National Institutional Ranking Framework (NIRF) in 2022, IIITB continues to excel in education and research, making it a preferred destination for aspiring technologists and future leaders.

AICTE | UGC

ABOUT upGrad

"upGrad is a leading global learning and workforce development company. We're on a single-minded mission of powering career success for every member of the global workforce as their trusted lifelong learning partner. Established in 2015, we have over 10 million learners who have upskilled in a range of online and offline programs from top universities in India and the world."

GSV EDTECH 150
Top 150 Most Transformative Growth Companies

BRANDON HALL GOLD AWARDS
15+ awards across 2 years

PROMISING BRAND 2022
The Economic Times

EDUCATION COMPANY OF THE YEAR
VC Circle, 2022

Entrepreneur India
EDUCATION INNOVATION AWARDS 2020
Best Communication & Teaching Platform

Ronnie Screwvala
Co-founder & Executive Chairman

Current Industry Trends

15 LPA

Average salary for Machine Learning Engineers in India (2025)

Source: Glassdoor

36.08%

CAGR of global Machine Learning market till 2030

Source: Statista

50%

Increase in search volume for ‘Data Science and Machine Learning’ since 2020

Source: Exploding Topics

74%

Annual increase in growth of AI and Machine Learning jobs since 2020

Source: Careerist

$503.4B

Size of global machine learning market by 2030

Source: Statista

Program Highlights

Here are the top reasons why you should consider this program

Future-Ready Curriculum
Master In-Demand and Trending Competencies

Personalised Learning Experience
Learning Experience Tailored to Your Needs

Specialisations
Specialise in two in-demand Data Science specialisations

In-Demand Tools
80+ Industry Tools, Languages, Libraries

Outcome-Driven Learning Experience
Personalised Portfolio-Building Support and Career Preparation Sessions

Best-in-Industry Experts
Decorated IIITB Faculty and Top Industry Practitioners

Golden Learning Ratio
Perfect Blend of Mathematics, Technology, and Business Understanding

Hands-on Learning
Solve 30+ Domain-Focused Assignments and Case Studies

Offline Graduation Function
On-Campus Graduation Ceremony for a Complete Program Experience

Program Impact Success Stories

Before upGrad

Tejasvy Gunturu

Application Development Associate

After upGrad

Application Development Analyst

Before upGrad

Vijeet Ved

System Engineer

After upGrad

IT Analyst

Before upGrad

Krish

Deputy Manager

After upGrad

Lead Engineer

Before upGrad

Vikas Kumar Gupta

Data Scientist

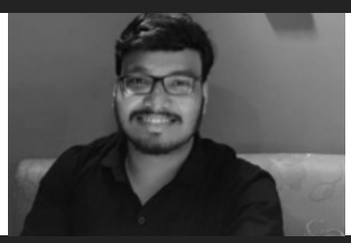

After upGrad

Senior Data Scientist

Before upGrad

Sunil Acharya

Software Engineer

After upGrad

AI Engineer

Our Transition Statistics

INR 1.23 Cr Highest Salary | 433% Highest Hike | 50% Average Hike

Program Completion Certificate

Earn valuable credentials with an Executive Diploma in Machine Learning and Artificial Intelligence. Join India's largest ML AI alumni network of over 10,000 professionals.

Earn Second Certificate

π Professional pack for dual specialisations

In just 3 more months learners can earn 1more specialization across both Data Science & Machine learning specializations

Faculty

Prof. Debabrata Das
Director, IIIT Bangalore

He has received his PhD from IIT-KGP.
His main areas of research are IoT and Wireless Access Network.

Prof. G. Srinivasaraghavan
IIIT Bangalore

Prof. Srinivasaraghavan has a PhD in Computer Science from IIT-K and 18 years of experience with Infosys Technologies and several other companies.

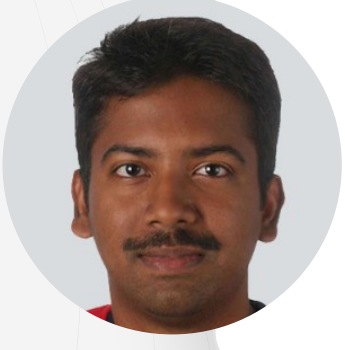

Prof. Dinesh Babu Jayagopi
IIIT Bangalore

Dr. Dinesh is currently an Associate Professor at IIIT-B where he heads the Multimodal Perception Lab. His research interests are in Audio-Visual Signal Processing, Machine Learning, and Social Computing. He obtained his doctorate from Ecole Polytechnic Federale Lausanne (EPFL), Switzerland.

Prof. Chandrashekar Ramanathan
Dean (Academics), IIIT Bangalore

Prof. Chandrashekar is a faculty member at IIIT-B since 2007 serving as professor, researcher and administrator. He has been working in the field of Computing for over 25 years in various capacities across industry and academia.

Prof. Tricha Anjali
Ex-Associate Dean, IIIT Bangalore

Prof. Anjali has a PhD from Georgia Institute of Technology as well as an integrated MTech (EE) from IIT Bombay.

Industry Experts

Abhishek Vijayvargia
Senior Data Scientist

Having worked with Microsoft as a Senior Data Scientist, he is an alumnus of IIT Kharagpur with 10+ years of experience in a Data Science domain

Microsoft
Ex-Senior Data Scientist

Anand
CEO

CEO, Gramener A gold medallist from IIM Bangalore, an alumnus of IIT Madras and London Business School, Anand is among the top 10 data scientists in India with 20 years of experience.

IIT MADRAS — Faculty
Infosys — Principal
BCG — Ex-Consultant

Manish Shukla
Head of Generative AI

Leading cutting-edge GenAI platform development at NatWest Group. Expertise in OpenAI products and MLOps for optimisation of operational efficiency and seamless project delivery with high user satisfaction.

Microsoft — Release Manager
NOKIA — Release Manager
Scrum Alliance — Certified Scrum Master

Deependra Singh
VP & Head of Data Science

Over 15 years of experience in leading analytics practices, data science, deep learning, and AI product development. Successfully led teams at Junglee Games, American Express Digital Business, and National Insurance Company, pioneering key projects like the analytics engine for the GOI PMJAY policy. Respected speaker at top educational institutes like IMT Hyderabad, BIT Mesra, and NMIMS.

Sajan Kedia
Sr. Engineering manager

Senior Engineering Manager, Hotstar Sajan has extensive experience in the field of ML, Big Data, Data Science, and AI.

Sr. Engineering manager
Machine Learning Engineer
Machine Learning Research Engineer

Mirza Rahim Baig
Startup Mentor

Analytics Lead, Zalando Mirza is a veteran professional with 10+ years of experience in application of data science, machine learning in e-commerce and healthcare.

Team Lead - Product
Marketing Analytics
Ex-Analytics Lead

Disclaimer: All product names, logos, brands, and trademarks are property of their respective owners. All company, product, and service names used on this platform are for identification and educational purposes only. The use of these names, logos, and brands does not imply any endorsement, affiliation, or partnership with the respective trademark holders.

Assignments and Case Studies from 12+ In-Demand Business Domains

Retail & Ecommerce
ETL Pipelining with Spark

Media & Entertainment
Data Analysis using SQL

Transportation
EDA
using Python

Education
Model Selection
using Sklearn

Civil Engineering
Classification using
CNNs

HR
Semantic Classification
using Word2Vec

Manufacturing
Regularisation using
Sklearn

Healthcare
Classification using
Sklearn

Law
RAG using
LangChain

InfoSec
Feature Engineering
using Sklearn

FMCG
Big Data Analysis
using Spark

BFSI
Sequence Data Prediction
using RNN

Learn by Doing

Your Program Journey

Phase 0

Math and Programming Bootcamp (12 weeks)

Phase I

Core Curriculum

(28 weeks, 15 credits)

Phase III

Capstone

(4 weeks, 7 credits)

Phase II

Specialisations

MLOps/GenAI

(22 weeks, 14 credits)

Applied Math and Programming Bootcamp

Personalise the initial 3 months of the program to your profile

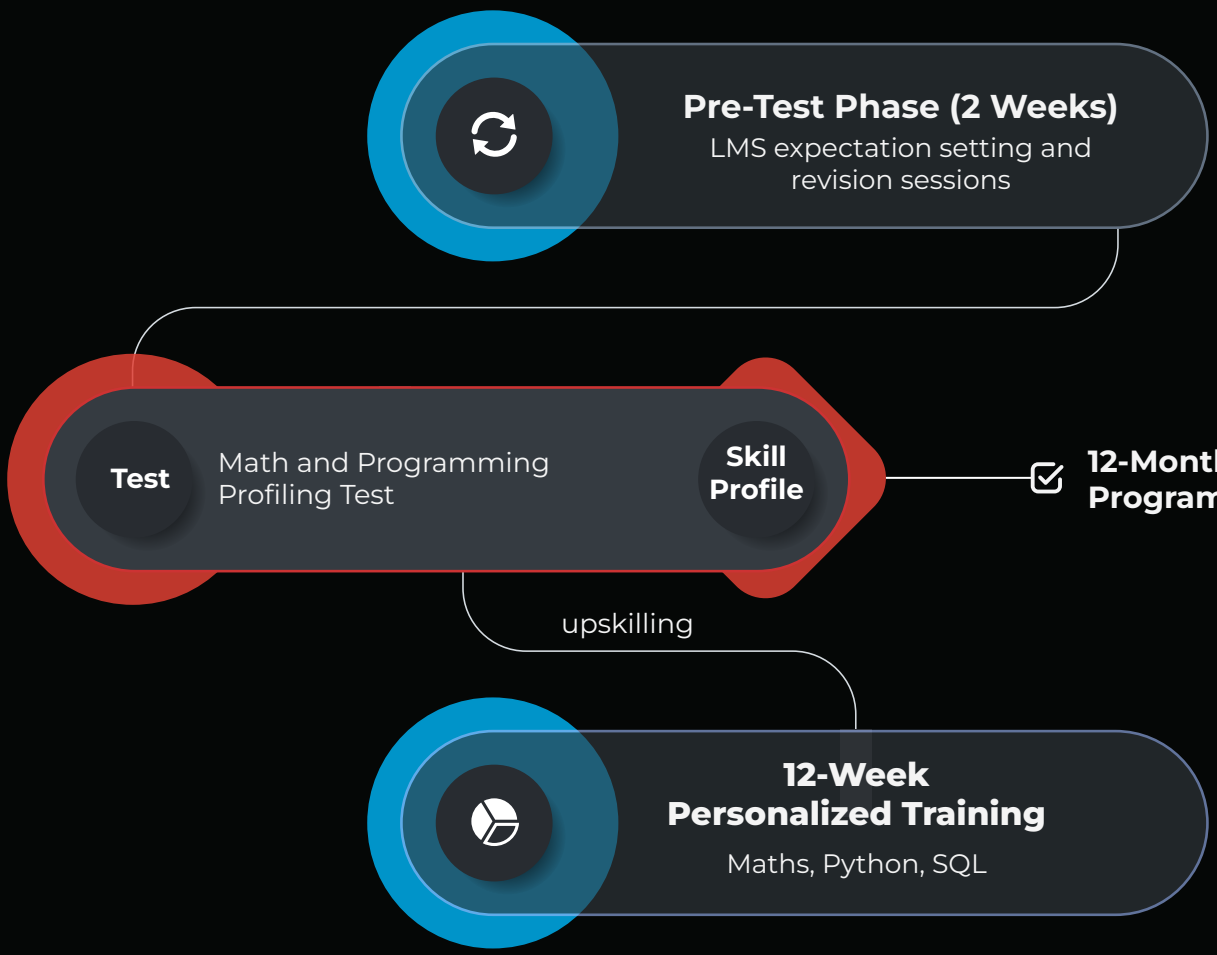

Topics: Sets, Combinatorics, Basics of Probability, Conditional Probability, Descriptive Statistics, Functions, Vector Algebra, Derivatives, Integrals, Coding Environments, Variables, Data Types, Syntax, Conditionals, Loops, Functions, Lists, Sets, Tuples, Dictionaries, Introduction to MySQL, Basic SQL Querying

Marks Structure: Total marks - 100

- Section A - **40 marks** (basic mathematics)
- Section B - **60 marks** (basic programming)
- Passing marks - 25 marks in section A & 35 marks in section B

No added cost to be paid for the bootcamp
We make sure that you are well-equipped to draw the most benefit from the program!

CORE CURRICULUM

Core Curriculum

The core phase of the curriculum will equip you with the most up-to-date and industry-relevant skills and technologies for data science and machine learning such as programming and mathematics, data analysis tools and techniques, cloud computing and big data analytics, and foundational topics in machine learning, deep learning, and natural language processing.

Topics

Advanced Mathematics for Data Science and Machine Learning

Master essential mathematical concepts to understand how to work with large amounts of data and train efficient machine learning models

- **Conditional Probability and Probability Distributions**
- **Advanced Linear Algebra and Linear Transformations**
- **Multivariate Calculus**

Advanced Programming for Data Science and Machine Learning

Wrangle real-world data using universal programming languages such as Python and SQL, and use GenAI for generating and debugging code faster

- **GenAI for Coding and Problem-Solving**
- **Object-Oriented Programming**
- **Python Data Science Libraries**
- **Database Design and SQL Querying with MySQL**
- **Introduction to NoSQL Databases**

Data Analysis and Exploration

Implement industry-standard statistical methods using tools such as Python, Tableau, and Power BI to analyse data and derive business insights

- **Data Analysis with Python**
- **Exploratory Data Analysis**
- **Inferential Statistics and Hypothesis Testing**
- **Data Analysis and Visualisation with Power BI and Tableau**

Cloud Computing and Big Data Fundamentals

Take your data processing and analysis workflows to the cloud and work with larger amounts of data to derive enterprise-scale business insights

- **Cloud Computing with AWS, GCP, Microsoft Azure**
- **Big Data Analysis with PySpark**

Foundations of Machine Learning

Train industry-standard machine learning models to automate insight generation and predict business metrics behaviour

- **Machine Learning Paradigms**
- **Linear and Logistic Regression**
- **K Nearest Neighbors**
- **Regularisation and Hyperparameter Tuning**
- **Decision Trees and Ensembles**
- **Clustering Models**

Deep Learning and Natural Language Processing

Build and train deep neural network models for different kinds of business data such as images and sequences

- **Artificial Neural Networks**
- **Convolutional and Recurrent Neural Networks**
- **Lexical, Syntactic, and Semantic Processing**

Deployment Fundamentals

Share and deploy your insights and machine learning models so that other collaborators can work with your contributions

- **Containerisation and Deployment Tools**
- **Version Control**

Projects

- **Querying with SQL**
  Analyse Spotify music data for targeted recommendations or NDAP insurance data for risk assessment
- **Exploratory Data Analysis**
  Analyse NYC taxi operations for efficient taxi positioning or US beer production data for better brewery operation management
- **Big Data Analysis**
  Analyse Mercari products data for better targeted recommendations or customer interaction data to enhance customer engagement
- **Linear Regression**
  Predict household energy consumption using appliance energy readings data to increase power consumption efficiency or parcel delivery time for Porter using historical delivery data for better planning and management
- **Deep Learning**
  Predict stock prices of Microsoft, Amazon, Google, IBM, using their historical stock price variations or temperature/pressure readings in Morocco using historical weather data

Tools

Disclaimer: All product names, logos, brands, and trademarks are property of their respective owners. All company, product, and service names used on this platform are for identification and educational purposes only. The use of these names, logos, and brands does not imply any endorsement, affiliation, or partnership with the respective trademark holders.

MLOps SPECIALISATION

MLOps Specialisation

The machine learning operations (MLOps) specialisation of the curriculum will equip you with core in-demand and industry relevant skills and technologies essential for ML engineers such as advanced machine learning methods, modern deep learning architectures, real-time data processing and end-to-end data pipeline creation and monitoring, and model pipelining and monitoring at scale.

Topics

Advanced Machine Learning

Train advanced industry-oriented machine learning models for enhanced predictive power and stronger business insight generation

- **Support Vector Machines and Naive Bayes**
- **Feature Engineering and Model Selection**
- **Dimensionality Reduction**
- **Time Series Analysis**
- **Association Rule Mining and Recommendation Systems**
- **Explainable AI**

Advanced Deep Learning and Generative AI

Design and train advanced industry-standard deep learning architectures, and master core AI principles such as attention mechanisms, transformers, and prompt engineering

- **Advanced CNN Architectures**
- **LSTMs and GRUs**
- **Transfer Learning Techniques**
- **Encoder-Decoder Architectures and Seq2Seq**
- **Machine Translation**
- **Attention Mechanisms and Transformers**
- **Fundamentals of Generative AI and Prompt Engineering**
- **Computer Vision, Variational Autoencoders, Generative Adversarial Networks**
- **Data and Model Security Principles**

Large-Scale Data Pipelining

Build complete end-to-end data pipelines and automate them to generate both batch-wise and real-time business insights

- **End-to-End Data Pipelining Fundamentals**
- **Pipeline Automation with AWS Lambda, GCP Functions, and Azure Automation**
- **Data Monitoring with Amazon CloudWatch, Google Cloud Monitoring, and Azure Monitor**
- **Feature Stores and Vector Databases**
- **Real-Time Analytics with Flink, Kafka, and Spark Streaming**
- **Real-Time Analytics with Amazon Kinesis, Google Cloud Pub/Sub and DataFlow, Azure Stream Analytics and Event Hubs**
- **Multicloud and Hybrid Cloud Operating Principles**

Machine Learning Model Pipelining

Build end-to-end industry-ready ML model pipelines and design their functional behaviour such as training and inference

- **Model Pipelining Principles**
- **Scheduling and Triggers**
- **Parallel Model Training and Real-Time Model Serving**
- **Data and Model Versioning**
- **Model Monitoring and System Design**

Projects

- **Feature Engineering and Model Selection**
  Predict fraudulent insurance claims using the Mendeley farmers insurance claims dataset or network intrusion events using historical network activity data
- **Semantic Classification**
  Fake News Detection, Job Role Classification
- **Real-Time Data Analytics: Develop a**
  real-time analytics pipeline for ecommerce data to enhance customer experience or a real-time patient health monitoring system for faster corrective actioning
- **Simulate and Retrigger Model Training Pipeline**

Tools

Disclaimer: All product names, logos, brands, and trademarks are property of their respective owners. All company, product, and service names used on this platform are for identification and educational purposes only. The use of these names, logos, and brands does not imply any endorsement, affiliation, or partnership with the respective trademark holders.

MLOps SPECIALISATION

Generative AI Specialisation

The generative artificial intelligence (GenAI) specialisation of the curriculum will equip you modern AI technologies and methods, particularly generative AI technologies, essential to data scientists and AI specialists, such as advanced machine learning methods, modern deep learning architectures, advanced prompt engineering and generative AI system design, information retrieval and retrieval-augmented generation, large language model (LLM) deployment, advanced computer vision and 3D vision, GenAI optimisations, and AI ethics.

Topics

Advanced Machine Learning

Train advanced industry-oriented machine learning models for enhanced predictive power and stronger business insight generation

- **Support Vector Machines and Naive Bayes**
- **Feature Engineering and Model Selection**
- **Dimensionality Reduction**
- **Time Series Analysis**
- **Association Rule Mining and Recommendation Systems**
- **Explainable AI**

Advanced Deep Learning for Generative AI

Design and train advanced industry-standard deep learning architectures, and master core AI principles such as attention mechanisms, transformers, and prompt engineering

- **Advanced CNN Architectures**
- **LSTMs and GRUs**
- **Transfer Learning Techniques**
- **Encoder-Decoder Architectures and Seq2Seq**
- **Machine Translation**
- **Attention Mechanisms and Transformers**
- **Fundamentals of Generative AI and Prompt Engineering**
- **Computer Vision, Variational Autoencoders, Generative Adversarial Networks**
- **Data and Model Security Principles**

GenAI System Design

Design and orchestrate generative AI systems to leverage the power of generative AI models and transform business operations

- **Advanced Prompt Engineering and GenAI System Design**
- **Prompting Multimodal Models**
- **LLM Frameworks such as LangChain and LLaMa Index**
- **LLM Evaluation Methods**
- **Data Security and Governance**
- **AI Ethics**

Advanced Generative AI

Develop AI-based cutting-edge industry-level systems for greater business efficiency such as retrieval-augmented generation (RAG) systems and multimodal GenAI model prompt engineering

- **Information Retrieval Principles**
- **Embeddings and Vector Databases**
- **RAG Architectures**
- **Agentic Systems and Multi-Agent Systems**
- **Advanced Multimodal GenAI Models**
- **LLM Deployment**
- **Advanced Computer Vision and 3D Vision**
- **GenAI Optimisations**

Projects

- **Feature Engineering and Model Selection**
  Predict fraudulent insurance claims using the Mendeley farmers insurance claims dataset or network intrusion events using historical network activity data

- **Semantic Classification**
  Fake News Detection, Job Role Classification

- **GenAI System Design**
  Analyse Amazon customer reviews to identify prevalent sentiments and themes to improve product offerings and enhance customer satisfaction or ChatGPT customer feedback to derive actionable insights for business improvement

- **RAG**
  Develop an RAG system to transform Long Beach County Municipal meetings transcripts into actionable insights for better organisational communication and decision making or historical legal documentations to optimise legal workflows

Tools

Disclaimer: All product names, logos, brands, and trademarks are property of their respective owners. All company, product, and service names used on this platform are for identification and educational purposes only. The use of these names, logos, and brands does not imply any endorsement, affiliation, or partnership with the respective trademark holders.

CAPSTONE PROJECTS

Capstone Projects

Capstone that Adapts to Your Preference

Infuse our Capstone with Your Data

Modify existing projects as per your industry data and problems

Bring Your Own Capstone

Work on a completely novel project of your choice and solve problems that excite you

Pre-Designed Industry Capstone

Choose one of our existing projects that cover in-demand trending industry domains

Bring Your Own Capstone

Design your own capstone project relevant to your domain and interest, and get feedback throughout your capstone stages

Identify a real-world problem relevant to your domain

Source datasets aligned with your business problem

Design and implement your solution

Document your efforts and present your findings

Continuous expert feedback at every step of capstone

Build A Strong Portfolio

GitHub

- **Commits**
  Demonstrate consistency, collaboration, and coding discipline
- **Code**
  Showcase well-documented repositories
- **Projects**
  Host end-to-end DS/ML/AI projects that highlight real-world problem-solving

GitHub helps with

- Validating coding skills
- Showing growth and consistency
- Being interview-ready for Tech roles

kaggle

- **Kernels**
  Highlight data processing and EDA methodologies
- **Ranking**
  Evaluate and reflect global standing among data science practitioners
- **Competitions**
  Demonstrate problem-solving under tight constraints

Kaggle helps with

- Building credibility in data science circles
- Applying learning to real datasets
- Speaking confidently in Tech interviews

LinkedIn

- **Headline**
  Concise summary of goals, competencies, and professional identity
- **Summary**
  Engaging overview of learnin and career journey
- **Projects**
  Showcase practical experience, outcomes, and skill application

LinkedIn helps with

- Improving visibility with recruiters
- Positioning better for job openings
- Networking with peers and mentors in the field

Microsoft-Certified Advantage –

This program doesn't just prepare you for the world of Data Science & AIML—it gives you the Microsoft edge.

Learners earn industry backed certification from upGrad in association with Microsoft by completing specially designed modules integrated into the program, boosting both credibility and career readiness.

Microsoft Learn content modules Certification Modules:

- Introduction to Generative AI Concepts
- Introduction to GitHub Copilot
- Design & Manage Analytics Solutions using Power BI
- Designing & Implementing a Data Science Solution on Azure

Rich and Dedicated Live Support

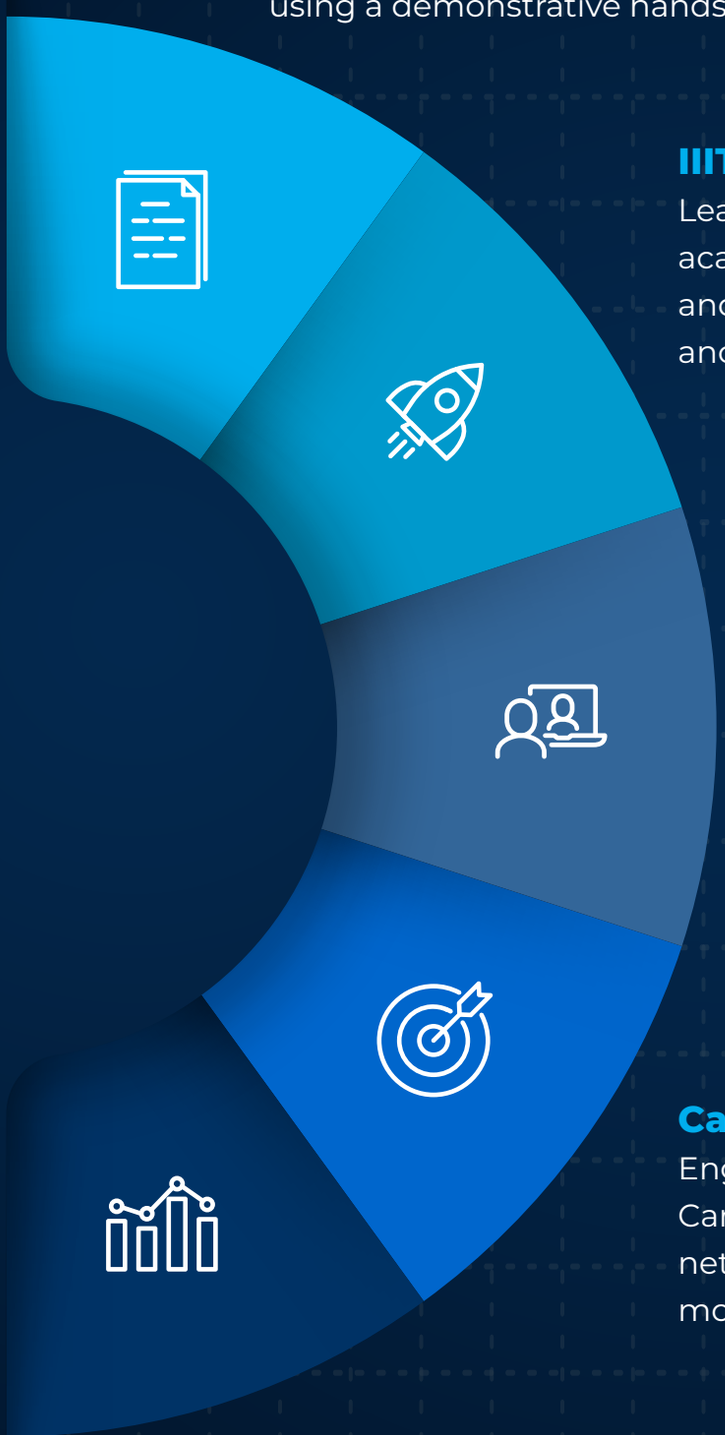

Industry Expert Sessions
Engage with industry practitioners as they help you master in-demand skills and concepts using a demonstrative hands-on approach

IIITB Faculty Sessions
Learn from some of the most accomplished academicians as they take your knowledge and understanding of data science to another level

Just-In-Time Interview Support
Participate in Technical and HR mock interviews designed to boost your confidence and prepare you to ace interviews.

Career Coaching Sessions
Engage in Career Coaching Sessions via Career preparation modules, High-impact networking events and Just-in-time mock interviews

Daily doubt resolution sessions
Join doubt resolution session slots, that are available daily, and have an expert available to resolve your queries for a smooth learning journey

Effective and Complete Career Support

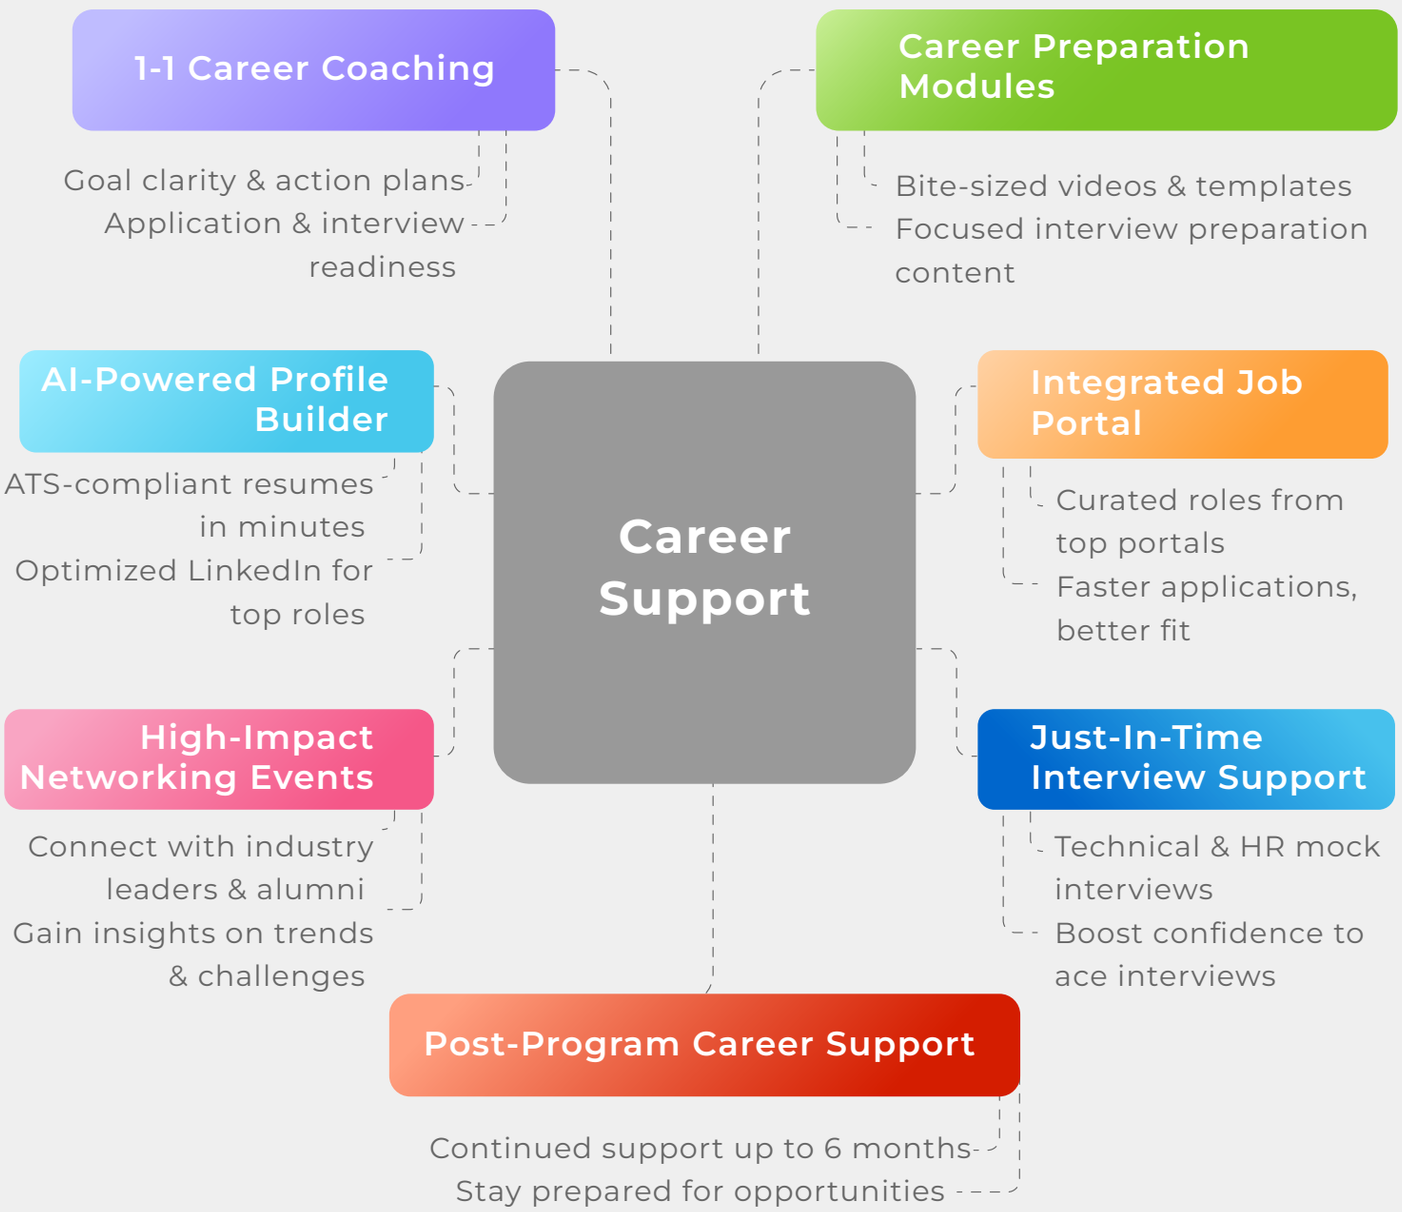

Student Support

Telegram channel for learner communications

Cohort Telegram channel for timely program updates and announcements.

Non-academic and non-technical query assistance

24 X 7 m-AI-ask bot support for non academic assistance.

Completion Support

Personalised assistance for smooth program completion, managing backlogs, and cohort deferrals with free and paid waiver options.

Financial benefits

Access benefits like referrals and repeats by sharing details with your program co-ordinator or student support executive.

Meet the Class

Industries Our Students Come From

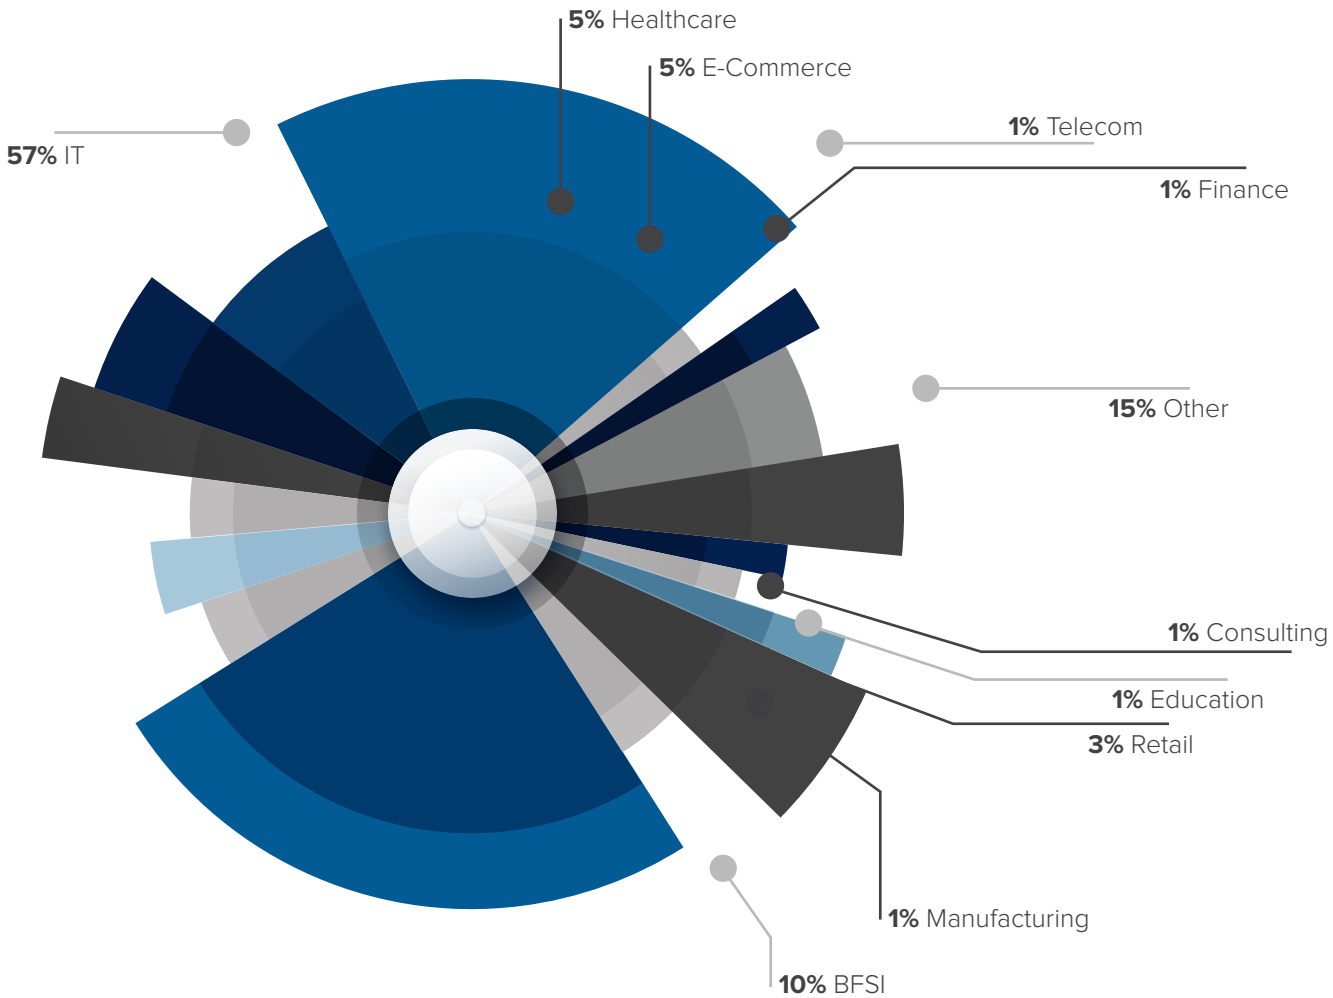

Work Experience

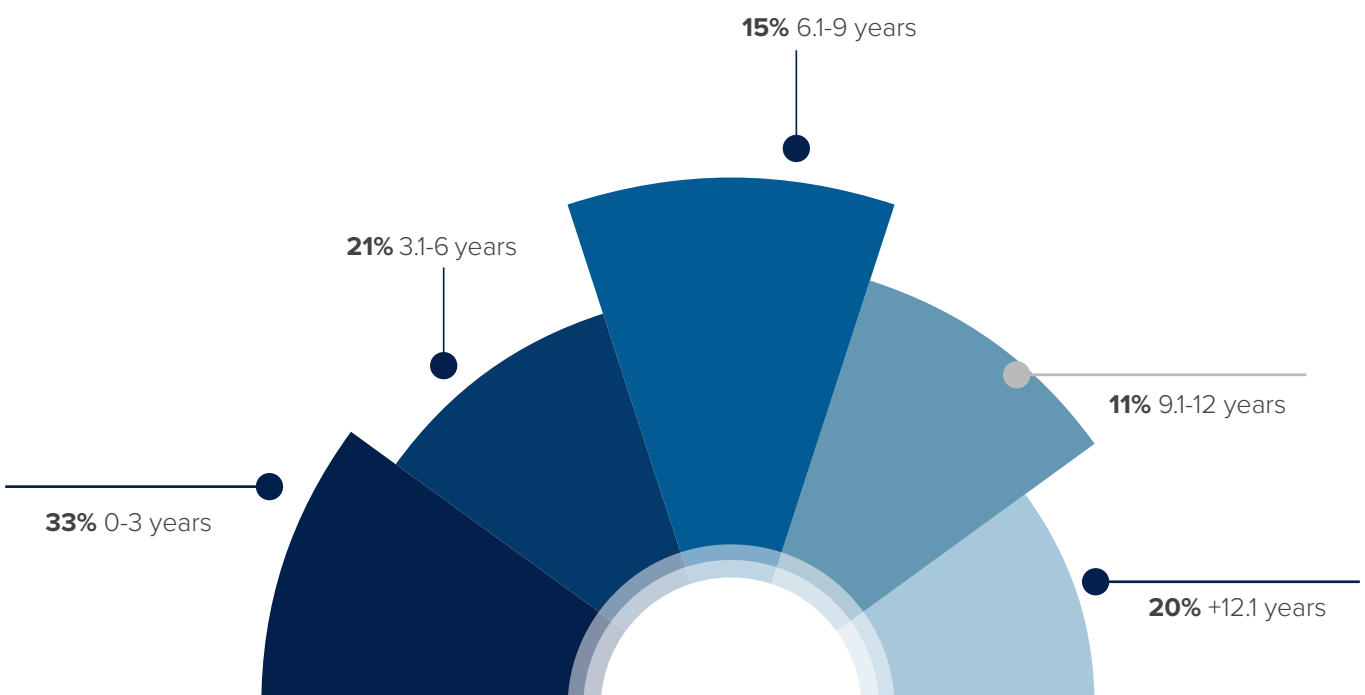

Professional pack

A single specialty is no longer enough.

For years, the T-shaped professional was considered the gold standard—someone with broad knowledge across multiple areas (the top of the “T”) and deep expertise in one core area (the vertical stem). This model worked well in stable environments where projects moved slowly, and roles were clearly divided

But today’s AI-driven, high-velocity world has exposed the limits of the T-shape. Knowledge silos, handoff delays, and rising complexity mean one deep skill is no longer enough

Today's employers need π-shaped professionals - someone with two deep, complementary specialisations while maintaining broad knowledge across domains, who can own entire project lifecycles without team dependencies, bridging the gap between development, operations, and analytics seamlessly

Introducing The π Professional Pack

Add an extra specialisation in just 3 months

STEP 1

Choose your primary foundation

Select your core specialisation:

Data Engineering, Data Analytics, Generative AI or MLOps

STEP 2

Add your strategic second specialty

In the latter half of your program, unlock a complete second track featuring:

- **Full curriculum coverage**
- **Live interactive sessions with industry experts**
- **Dedicated faculty support**
- **Expert mentorship across both domains**

Why is this a strategic career investment?

- **Expert in two complete specialisations**
- **Industry-validated skill combinations based on hiring data**
- **Hands-on experience in both domains**
- **Portfolio projects demonstrating π-shaped expertise**

Your Competitive Edge:

While others graduate with single specialisations, you emerge as a π-shaped leader ready to:

- **Own complete project lifecycles**
- **Bridge team silos and communication gaps**
- **Command premium positions in India's GCC & IT sectors**
- **Adapt quickly as AI continues reshaping the industry**

Popular specialisation combos

GenAI + MLOps

Why Gen AI + MLOps?

Enables end-to-end ownership of Gen AI models from fine-tuning to deployment, monitoring, and safety. This pairing creates a professional who can not only innovate but also reliably deliver and manage that innovation.

GenAI + Analytics

Why GenAI + Analytics?

Unlocks AI-driven business insights with a combination of natural language queries, auto-reporting, and data storytelling. This pairing is about creating smarter and more accessible business insights.

MLOps + Data Engineering

Why MLOps + Data Engineering?

This is the ultimate combination for building robust and scalable AI infrastructure. Enables you to build an automated, end-to-end machine learning factory that is efficient, reliable, and ready for production demands.

GenAI + Data Engineering

Why GenAI + Data Engineering?

Let's combine Gen AI innovation with solid data pipelines scalable AI products. This pairing ensures that innovative GenAI projects are built on a solid foundation, allowing them to be scalable and successful in the long run.

LIVERPOOL JOHN MOORES UNIVERSITY

Option to articulate to a Master's degree from Liverpool John Moores University after successful completion of the program

LJMU

Enrol in 4 small steps, Then take a giant leap.

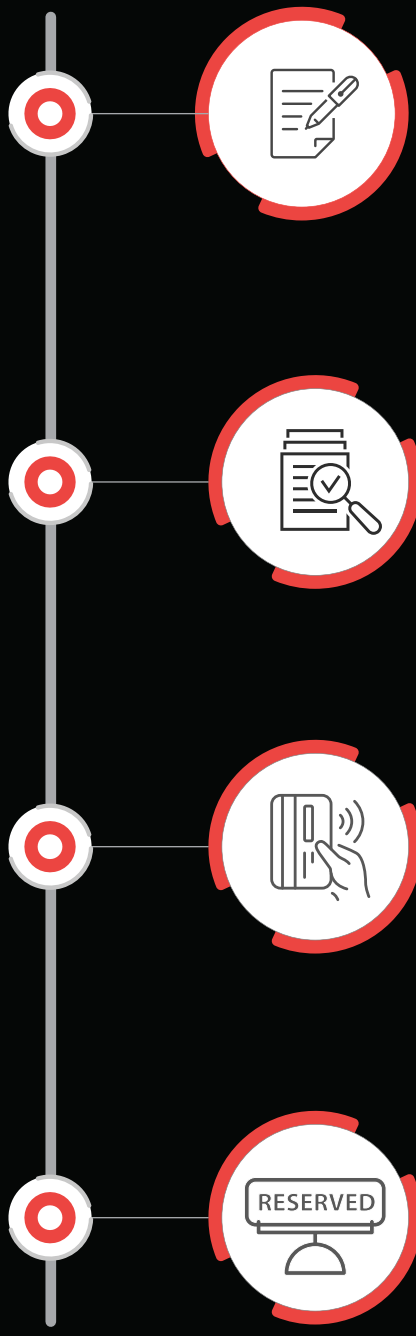

STEP 1

Submit Your Application

STEP 2

Get Shortlisted

STEP 3

Receive Your Offer

STEP 4

Block Your Seat

Eligibility Criteria

Bachelor's or Master's Degree or its equivalent in any discipline with minimum 50% aggregate mark or equivalent CGPA.

Write to us:
admissions@upgrad.com

Visit us at:
www.upgrad.com

Call: 1800-210-2020

upGrad

LET'S TALK