Interview Process (3 Stage)

Stage 01: General Interview

- An initial discussion to assess your background, overall suitability, and communication skills.
- Strong preparation is expected as this stage determines your eligibility to move forward.
- Please be prepared to discuss specific commercial projects you have built for clients after earning your degree in a relevant field (preferably Software Engineering Specializing AI / Data Science).
- Avoid vague or generic responses — detailed explanations of your practical contributions, architectural decisions, challenges faced, and how you solved them will be critical for success.

Stage 02: Technical Interview

An in-depth conversation on fundamental and advanced Machine Learning concepts, troubleshooting, research skills, and MLOps practices. This stage will cover:

- Choosing the right approach for a given problem.
- Optimising existing models for performance, scale, and accuracy.
- Researching and prototyping when existing libraries/models are insufficient.
- Deploying models reliably into production systems.

This document is strictly private and confidential. It is intended only for the individual(s) addressed and must not be shared, copied, or used by anyone other than the designated recipient without prior written authorization from Intelligent Mind Labs.

Your prior experience in building and delivering commercial AI products will be highly relevant here. The session will be recorded and further reviewed by a panel of AI experts of the respective fields.

Bonus Criteria (A strong mathematical foundation will be a significant advantage):

- Solid understanding of linear algebra, calculus, probability, and statistics (sufficient to derive or adapt models).
- Knowledge of optimization techniques (e.g., gradient descent, convex vs. non-convex problems).
- Familiarity with the bias–variance tradeoff and concepts of overfitting/underfitting.
- Experience with evaluation metrics across tasks such as classification, regression, ranking, and clustering.

Stage 03: Technical Assignment

- Candidates who successfully clear the technical interview will be given a hands-on assignment.
- This assignment is designed to evaluate your ability to apply basic to advance ML concepts to real-world problems. Upon completion, you will advance to a comprehensive face-to-face interview to review your work.
- Use of GenAI in this assignment is **strictly prohibited.**

This document is strictly private and confidential. It is intended only for the individual(s) addressed and must not be shared, copied, or used by anyone other than the designated recipient without prior written authorization from Intelligent Mind Labs.

A hands-on assignment covering end-to-end ML pipelines:

- Data Pipeline
- Model Training Pipeline
- Validation & Evaluation Pipeline
- Model Packaging & Deployment Pipeline
- Monitoring & Maintenance Pipeline
- MLOps Orchestration Pipeline (Advance use of Docker is essential in this stage)

This document is strictly private and confidential. It is intended only for the individual(s) addressed and must not be shared, copied, or used by anyone other than the designated recipient without prior written authorization from Intelligent Mind Labs.