

Training Course on A/B Testing and Experimentation for ML Models

Professional Development Training Course Outline

Category	Data Science	Duration	10 Days
Base Fee	\$3,000 USD	Delivery Mode	Online / On-site Hybrid

Course Introduction

This comprehensive training course dives deep into the critical domain of **A/B testing and online experimentation** specifically tailored for **Machine Learning (ML) models**. In today's data-driven world, optimizing ML model performance in real-world scenarios is paramount for **business growth** and **user engagement**. Training Course on A/B Testing & Experimentation for ML Models: Designing and Analyzing Online Experiments equips participants with the **statistical rigor**, **experimental design principles**, and **practical tools** necessary to confidently launch, analyze, and iterate on ML model deployments, moving beyond offline metrics to **live production validation**.

Participants will learn to design robust experiments, interpret results with **statistical significance**, and make **data-backed decisions** for continuous ML model improvement. The course emphasizes both the theoretical foundations and hands-on application, covering everything from formulating **testable hypotheses** to advanced techniques like **multi-armed bandits** and **causal inference**. Master the art of **controlled experiments** to unlock superior model performance, drive **conversion rate optimization (CRO)**, and deliver exceptional **user experiences** in dynamic online environments.

Programme Curriculum

Training Course on A/B Testing & Experimentation for ML Models: Designing and Analyzing Online Experiments

Introduction

This comprehensive training course dives deep into the critical domain of **A/B testing and online experimentation** specifically tailored for **Machine Learning (ML) models**. In today's data-driven world, optimizing ML model performance in real-world scenarios is paramount for **business growth** and **user engagement**. Training Course on A/B Testing & Experimentation for ML Models: Designing and Analyzing Online Experiments equips participants with the **statistical rigor**

, **experimental design principles**, and **practical tools** necessary to confidently launch, analyze, and iterate on ML model deployments, moving beyond offline metrics to **live production validation**.

Participants will learn to design robust experiments, interpret results with **statistical significance**, and make **data-backed decisions** for continuous ML model improvement. The course emphasizes both the theoretical foundations and hands-on application, covering everything from formulating **testable hypotheses** to advanced techniques like **multi-armed bandits** and **causal inference**.

Master the art of **controlled experiments** to unlock superior model performance, drive **conversion rate optimization (CRO)**, and deliver exceptional **user experiences** in dynamic online environments.

Course Duration

10 days

Course Objectives

1. Grasp core concepts of online experimentation and their application in machine learning workflows.
2. Learn to structure A/B, A/A, and multivariate tests for ML models, ensuring **validity** and **statistical power**.
3. Develop clear, measurable hypotheses specific to ML model performance and business objectives.
4. Identify and define key evaluation metrics for online experiments.
5. Determine statistically significant sample sizes and optimal experiment durations for reliable results.
6. Gain practical skills in utilizing industry-standard A/B testing tools and platforms for ML deployments.
7. Apply **hypothesis testing**, **p-values**, and **confidence intervals** to interpret A/B test outcomes.
8. Differentiate between statistically significant findings and their real-world business impact.
9. Account for user behavior changes over time in long-running experiments.
10. Delve into **multi-armed bandits**, **sequential testing**, and **causal inference** for adaptive optimization.
11. Identify common pitfalls and biases in online experiments and implement corrective measures.
12. Establish a continuous feedback loop for data-driven model improvement and deployment.
13. Effectively present A/B test findings and recommendations to technical and non-technical stakeholders.

Organizational Benefits

- Faster, data-driven decisions on model efficacy lead to quicker time-to-market for improved features.
- Validate model changes in a controlled environment before full rollout, minimizing negative impacts on user experience or revenue.
- Continuously test and refine ML-driven features based on actual user behavior, leading to more personalized and satisfying interactions.
- Identify and scale winning model versions that directly contribute to key business metrics.
- Foster a culture of experimentation and evidence-based innovation across product, engineering, and data science teams.
- Direct development efforts towards changes proven to have a positive impact, avoiding wasted resources on ineffective ideas.

- Stay ahead by consistently delivering superior ML-powered products and services through continuous optimization.

Target Audience

1. Machine Learning Engineers
2. Data Scientists
3. Product Managers
4. AI/ML Researchers
5. Software Engineers
6. Analytics Professionals
7. Growth Marketers
8. Technical Leads and Team Managers

Course Outline

Module 1: Introduction to A/B Testing for ML

- What is A/B testing and why is it crucial for ML models?
- Distinction between offline model evaluation and online experimentation.
- Benefits of A/B testing in reducing deployment risk and driving growth.
- Ethical considerations in A/B testing and responsible experimentation.
- The experimentation lifecycle: from hypothesis to decision.
- **Case Study:** Netflix's personalized recommendations (testing new algorithms).

Module 2: Statistical Foundations for Experimentation

- Review of core statistical concepts: mean, variance, standard deviation.
- Understanding probability distributions: Normal, Bernoulli, Binomial.
- Central Limit Theorem and its importance in A/B testing.
- Introduction to hypothesis testing: Null and Alternative Hypotheses (H_0 ?, H_1 ?).
- Type I and Type II errors: Alpha (?) and Beta (?), and their trade-offs.
- **Case Study:** Google Ads optimizing click-through rates (CTRs) via small statistical uplifts.

Module 3: Designing Robust Experiments

- Key elements of experimental design: control group, treatment group, randomization.
- Common experiment types: A/A, A/B, A/B/n, Multivariate Testing.
- Choosing the right experimental unit: users, sessions, events.
- Avoiding common design pitfalls: selection bias, novelty effect.
- Segmentation and targeting in A/B tests for personalized experiences.
- **Case Study:** Amazon testing new product page layouts to increase conversion.

Module 4: Defining Metrics and KPIs for ML Experiments

- North Star metrics vs. Guardrail metrics.
- Primary and secondary metrics for ML model evaluation (e.g., accuracy, precision, recall, latency, revenue per user, conversion rate).
- Metric sensitivity and power.
- Proxy metrics and their limitations.
- Defining Overall Evaluation Criteria (OEC).
- **Case Study:** Spotify experimenting with different recommendation algorithms and measuring listen time and user retention.

Module 5: Sample Size Calculation and Experiment Duration

- Understanding statistical power and its relationship to sample size.
- Calculating minimum detectable effect (MDE).
- Formulas for sample size calculation for different metric types (proportions, means).
- Practical considerations for experiment duration: traffic, seasonality, ramp-up time.
- Using A/B test calculators and statistical software.
- **Case Study:** Facebook's early A/B tests for feature rollouts with limited user segments.

Module 6: Running Experiments: Practical Implementation

- Setting up experimentation platforms (e.g., Optimizely, Split.io, custom solutions).
- Feature flagging and remote configuration for ML models.
- Data collection best practices for online experiments.
- Monitoring experiment health and detecting anomalies.
- Infrastructure considerations for scalable experimentation.
- **Case Study:** LinkedIn optimizing connection recommendations by dynamically allocating users to different model versions.

Module 7: Analyzing Experiment Results: Statistical Significance

- Performing t-tests, chi-squared tests, and ANOVA for A/B test analysis.
- Interpreting p-values and confidence intervals.
- Multiple comparisons problem and correction techniques (Bonferroni, FDR).
- Practical vs. statistical significance: understanding business impact.
- Reporting and visualizing A/B test results effectively.
- **Case Study:** Booking.com continuously testing website elements and measuring booking conversion rates.

Module 8: Advanced Topics in Experimentation

- Sequential testing and early stopping rules.
- Bayesian A/B testing: advantages and applications.
- Switchback experiments and their use cases.
- Network effects and interference in experiments.
- Causal inference techniques for understanding "why."
- **Case Study:** Uber's experimentation platform using advanced techniques to optimize ride-hailing dynamics.

Module 9: Debugging and Troubleshooting Experiments

- Common sources of A/B test invalidity: biased assignment, data pollution.
- Pre-test checks and A/A tests for validation.
- Monitoring experiment data quality and integrity.
- Identifying and resolving technical issues during experiments.
- Post-experiment analysis and root cause identification for unexpected results.
- **Case Study:** A retail e-commerce company detecting data logging errors that skewed experiment results.

Module 10: Experimentation for ML Model Iteration

- Integrating A/B testing into the ML development lifecycle (MLOps).
- Continuous deployment and continuous experimentation for models.
- A/B testing for model selection, hyperparameter tuning, and feature engineering.
- Evaluating fairness and bias in ML models through A/B tests.
- A/B testing for model retraining strategies.

- **Case Study:** Google Search ranking algorithm updates and their validation through large-scale online experiments.

Module 11: Multi-Armed Bandits (MAB) for Dynamic Optimization

- Introduction to Multi-Armed Bandits: Explore vs. Exploit dilemma.
- Common MAB algorithms: Epsilon-Greedy, Upper Confidence Bound (UCB), Thompson Sampling.
- When to use MABs instead of traditional A/B tests.
- Applications of MABs in personalized content, recommendations, and advertising.
- Implementing and monitoring MAB systems.
- **Case Study:** News websites dynamically optimizing headlines using MABs to maximize engagement.

Module 12: Causal Inference in A/B Testing

- Understanding causality vs. correlation in experimentation.
- Potential outcomes framework and treatment effects.
- Addressing confounding variables and bias in observational studies.
- Regression discontinuity and difference-in-differences for quasi-experiments.
- Advanced causal inference methods (e.g., propensity score matching) for non-experimental data.
- **Case Study:** A ride-sharing company analyzing the causal impact of a new pricing model on rider retention.

Module 13: Building an Experimentation Culture

- The importance of organizational buy-in for experimentation.
- Establishing clear processes and best practices for A/B testing.
- Promoting a data-driven mindset across teams.
- Measuring the ROI of experimentation.
- Sharing learnings and fostering continuous improvement.
- **Case Study:** Microsoft's journey in building a company-wide experimentation platform and culture.

Module 14: A/B Testing Beyond Core Model Performance

- Testing ML model explainability and interpretability improvements.
- A/B testing for model monitoring and alerting systems.
- Evaluating the user interface (UI) and user experience (UX) of ML-powered features.
- A/B testing for model deployment strategies (e.g., canary rollouts).
- Security and privacy considerations in online experimentation.
- **Case Study:** An online streaming service A/B testing different ways to explain why a movie was recommended.

Module 15: Future Trends and Best Practices

- AI-powered A/B testing and automation.
- Personalized experimentation and adaptive learning.
- Federated learning and privacy-preserving experimentation.
- Ethical AI and responsible experimentation guidelines.
- Emerging tools and technologies in the experimentation landscape.
- **Case Study:** Discussing future challenges and opportunities for experimentation in a rapidly evolving ML landscape.

Training Methodology

This course employs a **blended learning approach** combining:

- **Interactive Lectures:** Core concepts explained with clear examples and visual aids.
- **Hands-on Labs & Coding Exercises:** Practical application using Python (Pandas, NumPy, SciPy, Statsmodels) and popular A/B testing simulation libraries.
- **Real-world Case Studies:** In-depth analysis of successful and challenging A/B tests from leading tech companies.
- **Group Discussions & Collaborative Problem-Solving:** Encouraging peer learning and diverse perspectives.
- **Q&A Sessions:** Dedicated time for addressing participant queries and clarifying concepts.
- **Capstone Project:** Participants will design, simulate, and analyze an A/B test for an ML model scenario.

Register as a group from 3 participants for a Discount

Send us an email: info@fineskilltrainingcenter.com or call +254769199797

Certification

Upon successful completion of this training, participants will be issued with a globally- recognized certificate.

Tailor-Made Course

We also offer tailor-made courses based on your needs.

Key Notes

Upcoming Scheduled Sessions

Start Date	End Date	Location	Price
07 Sep 2026	18 Sep 2026	Physical	\$3,000
14 Sep 2026	25 Sep 2026	Physical	\$3,000
21 Sep 2026	02 Oct 2026	Physical	\$3,000
28 Sep 2026	09 Oct 2026	Physical	\$3,000
05 Oct 2026	16 Oct 2026	Physical	\$3,000
12 Oct 2026	23 Oct 2026	Physical	\$3,000
19 Oct 2026	30 Oct 2026	Physical	\$3,000

Start Date	End Date	Location	Price
26 Oct 2026	06 Nov 2026	Physical	\$3,000

Ready to enroll or request a customized program? Book online or contact us:

[Register Online Now](#)

Email: info@fineskilltrainingcenter.com | Website: www.fineskilltrainingcenter.com