

Training Course on Optimizing Machine Learning Algorithms (Advanced)

Professional Development Training Course Outline

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

Course Introduction

In the rapidly evolving landscape of Artificial Intelligence and Machine Learning, the ability to **optimize machine learning algorithms** is paramount for achieving cutting-edge performance and deploying robust, scalable AI solutions. This advanced training course dives deep into the intricate world of **gradient descent variants** and the complex **optimization landscapes** that govern model training. Participants will gain a sophisticated understanding of how to fine-tune algorithms, navigate challenging loss surfaces, and unlock the full potential of their **deep learning models** and predictive analytics.

Training Course on Optimizing Machine Learning Algorithms (Advanced) empowers data scientists, ML engineers, and researchers to transcend basic model training, equipping them with the **advanced optimization techniques** necessary to build highly efficient, accurate, and stable machine learning systems. Through a blend of theoretical foundations and practical, hands-on exercises, attendees will master the art of **hyperparameter tuning**, **convergence acceleration**, and mitigating common optimization challenges like **local minima** and **saddle points**, ultimately driving superior **model performance** and business impact.

Programme Curriculum

Training Course on Optimizing Machine Learning Algorithms (Advanced)

Introduction

In the rapidly evolving landscape of Artificial Intelligence and Machine Learning, the ability to **optimize machine learning algorithms** is paramount for achieving cutting-edge performance and deploying robust, scalable AI solutions. This advanced training course dives deep into the intricate world of **gradient descent variants** and the complex **optimization landscapes** that govern model training. Participants will gain a sophisticated understanding of how to fine-tune algorithms, navigate challenging loss surfaces, and unlock the full potential of their **deep learning models** and

predictive analytics.

Training Course on Optimizing Machine Learning Algorithms (Advanced) empowers data scientists, ML engineers, and researchers to transcend basic model training, equipping them with the **advanced optimization techniques** necessary to build highly efficient, accurate, and stable machine learning systems. Through a blend of theoretical foundations and practical, hands-on exercises, attendees will master the art of **hyperparameter tuning**, **convergence acceleration**, and mitigating common optimization challenges like **local minima** and **saddle points**, ultimately driving superior **model performance** and business impact.

Course Duration

10 days

Course Objectives

Upon completion of this intensive training, participants will be able to:

1. Comprehend and implement advanced gradient descent algorithms, including Stochastic Gradient Descent (SGD), Mini-Batch Gradient Descent, Adam, RMSprop, and Adagrad.
2. Analyze and interpret complex loss landscapes, identifying and addressing issues like local minima, saddle points, and plateaus.
3. Apply techniques for convergence acceleration, such as momentum, Nesterov accelerated gradient, and learning rate schedules.
4. Utilize advanced regularization techniques (L1, L2, Dropout, Batch Normalization) to prevent overfitting and enhance model generalization.
5. Employ sophisticated hyperparameter tuning methods like Grid Search, Random Search, Bayesian Optimization, and Automated Machine Learning (AutoML).
6. Explore the principles and applications of second-order optimization methods like Newton's method and Quasi-Newton methods (BFGS, L-BFGS).
7. Recognize and mitigate problems arising from ill-conditioned loss functions, including strategies for preconditioning.
8. Implement distributed optimization and parallel training techniques for efficient model training on massive datasets.
9. Critically assess and compare the performance of different optimization algorithms using appropriate evaluation metrics.
10. Develop systematic approaches to debug and troubleshoot common optimization challenges in real-world scenarios.
11. Optimize specific deep learning architectures (e.g., CNNs, RNNs, Transformers) for improved performance.
12. Deepen understanding of adaptive learning rate methods and their impact on training stability and speed.
13. Gain insights into emerging optimization research and cutting-edge algorithms in the ML community.

Organizational Benefits

- Develop and deploy ML models with significantly higher accuracy, robustness, and generalization capabilities, leading to better predictions and decisions.
- Accelerate the training and iteration process of complex ML models, reducing time-to-market for AI-powered products and services.
- Efficiently train models on available computational resources, leading to reduced infrastructure costs and energy consumption.

- Equip teams with the expertise to tackle challenging ML problems that require advanced optimization, such as large-scale **neural networks** and **reinforcement learning**.
- Build more reliable and generalizable models that perform consistently on unseen data, minimizing business risks associated with poor model performance.
- Foster a team of highly skilled ML professionals capable of implementing state-of-the-art **AI solutions**, distinguishing the organization in the market.
- Drive new opportunities and innovations by enabling the development of more sophisticated and performant **machine learning applications**.

Target Audience

1. Data Scientists
2. Machine Learning Engineers .
3. AI Researchers.
4. Deep Learning Practitioners
5. Software Engineers
6. Ph.D. Students and Academics in AI, Computer Science, and related fields.
7. Professionals involved in AI/ML Model Deployment and Performance Monitoring.
8. Anyone with a solid grasp of fundamental machine learning concepts and calculus seeking to enhance their optimization skills.

Course Outline

Module 1: Foundations of Optimization in ML

- Review of Convexity and Non-Convexity
- Understanding Loss Functions and Objective Surfaces
- The Role of Gradients in Optimization
- Introduction to First-Order vs. Second-Order Methods
- Challenges in High-Dimensional Optimization

Module 2: Batch and Stochastic Gradient Descent (SGD) Revisited

- Deep Dive into Batch Gradient Descent Mechanics
- The Power and Noise of Stochastic Gradient Descent
- Mini-Batch Gradient Descent: The Industry Standard
- Practical Considerations for Batch Size Selection
- **Case Study:** Optimizing a Logistic Regression model for credit risk prediction using different batch sizes.

Module 3: Momentum and Nesterov Accelerated Gradient

- Overcoming Oscillations with Momentum
- Nesterov Accelerated Gradient for Faster Convergence
- Implementing Momentum-based Optimizers from Scratch
- Visualizing Trajectories on Optimization Landscapes
- **Case Study:** Accelerating training of a feedforward neural network on image classification using momentum.

Module 4: Adaptive Learning Rate Methods - Adagrad & RMSprop

- Adaptive Learning Rates: Why They Matter
- Adagrad: Per-Parameter Learning Rates
- RMSprop: Addressing Adagrad's Diminishing Learning Rate

- Comparing Performance and Stability
- **Case Study:** Training a sentiment analysis model with word embeddings, comparing Adagrad and RMSprop for convergence.

Module 5: The Adam Optimizer and its Variants

- Adam: Combining Momentum and Adaptive Learning Rates
- Understanding Adam's Bias Correction Mechanism
- Exploring AdamW for Weight Decay Regularization
- Other Adam-like Optimizers (e.g., Nadam, AMSGrad)
- **Case Study:** Fine-tuning a pre-trained Transformer model for natural language understanding using AdamW.

Module 6: Optimization Landscapes and Challenges

- Visualizing Complex Loss Landscapes (e.g., Saddle Points, Local Minima)
- Understanding the Vanishing and Exploding Gradient Problem
- The Impact of Initialization on Optimization
- Techniques to Escape Local Minima and Saddle Points
- **Case Study:** Analyzing the loss landscape of a deep convolutional neural network and identifying strategies to overcome convergence issues.

Module 7: Regularization Techniques for Robust Optimization

- L1 and L2 Regularization (Weight Decay)
- Dropout: A Powerful Regularization Strategy
- Batch Normalization for Stabilizing Training
- Early Stopping and Data Augmentation as Regularizers
- **Case Study:** Improving the generalization of a medical image segmentation model by applying a combination of regularization techniques.

Module 8: Hyperparameter Tuning Strategies

- Manual Tuning vs. Automated Approaches
- Grid Search and Random Search: Practical Implementations
- Bayesian Optimization for Efficient Hyperparameter Discovery
- Introduction to AutoML Tools and Frameworks
- **Case Study:** Optimizing hyperparameters for a Gradient Boosting Machine (GBM) model to predict customer churn.

Module 9: Second-Order Optimization Methods

- Introduction to Newton's Method for Optimization
- Quasi-Newton Methods: BFGS and L-BFGS
- Approximating the Hessian Matrix
- Advantages and Limitations of Second-Order Methods
- **Case Study:** Applying L-BFGS to optimize a small-scale neural network or a maximum likelihood estimation problem.

Module 10: Advanced Topics in Learning Rate Schedules

- Cyclical Learning Rates
- Learning Rate Warm-up
- Cosine Annealing Learning Rate Schedule
- Dynamic Adjustment Strategies

- **Case Study:** Implementing and evaluating different learning rate schedules for training a large-scale image recognition model.

Module 11: Distributed and Parallel Optimization

- Data Parallelism vs. Model Parallelism
- Horovod and Distributed Training Frameworks
- Synchronous vs. Asynchronous Updates
- Challenges in Distributed Optimization (e.g., Communication Overhead)
- **Case Study:** Scaling the training of a deep learning model across multiple GPUs or machines for a big data analytics task.

Module 12: Debugging and Troubleshooting Optimization

- Common Signs of Optimization Issues (e.g., Divergence, Slow Convergence)
- Monitoring Loss Curves and Gradients
- Techniques for Identifying Root Causes of Problems
- Strategies for Remediation and Recovery
- **Case Study:** Diagnosing and resolving optimization challenges in a real-world project, such as a recommender system with noisy data.

Module 13: Optimization for Specific ML Architectures

- Optimizing Convolutional Neural Networks (CNNs) for Computer Vision
- Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) Optimization
- Optimization Challenges in Generative Adversarial Networks (GANs)
- Strategies for Reinforcement Learning Optimization
- **Case Study:** Enhancing the training stability and performance of a GAN for realistic image generation.

Module 14: Practical Considerations and Best Practices

- Choosing the Right Optimizer for Your Problem
- Reproducibility in Optimization
- The Importance of Data Preprocessing
- Ethical Considerations in Model Optimization
- Deployment Challenges Related to Optimization
- **Case Study:** Developing an end-to-end optimized ML pipeline for a real-time anomaly detection system.

Module 15: Emerging Trends in Optimization

- Beyond First and Second Order: Recent Research
- Meta-Learning for Optimization
- Neural Architecture Search (NAS) and Optimization
- Hardware-Aware Optimization
- Future Directions in ML Optimization
- **Case Study:** Discussing the potential impact of a recent research paper on a specific industry application, e.g., using meta-learning for faster adaptation in federated learning.

Training Methodology

This course employs a **blended learning approach** designed for maximum engagement and practical skill development:

- **Interactive Lectures:** In-depth theoretical discussions on gradient descent variants, optimization landscapes, and advanced techniques.
- **Hands-on Coding Labs:** Extensive practical sessions using Python with popular ML frameworks (TensorFlow, PyTorch, scikit-learn) for implementing and experimenting with various optimizers.
- **Real-World Case Studies:** Application of learned concepts to solve practical challenges across diverse domains (e.g., computer vision, NLP, finance).
- **Collaborative Exercises:** Group activities and discussions to foster peer learning and problem-solving.
- **Live Demonstrations:** Expert-led demonstrations of complex optimization scenarios and debugging techniques.
- **Code Review and Feedback:** Opportunities for participants to present their code and receive constructive feedback.
- **Q&A Sessions:** Dedicated time for addressing participant questions and clarifying concepts.
- **Practical Projects:** A culminating project where participants apply advanced optimization techniques to a real-world dataset.

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

- The participant must be conversant with English.
- Upon completion of training the participant will be issued with an Authorized Training Certificate
- Course duration is flexible and the contents can be modified to fit any number of days.

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

Start Date	End Date	Location	Price
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
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