Machine Learning Training with BigML
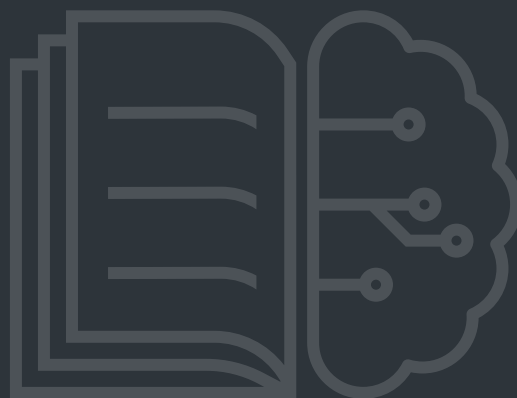
To Machine Learning From Scratch

This is a practical, hands-on course to introduce the basic concepts of Machine Learning. The course uses basic examples that the trainees can follow step-by-step to implement their own models. The ultimate goal is to teach students which Machine Learning tasks should be used to solve particular problems or use cases. Students will also learn about the steps and workflows to build real-world Machine Learning applications end-to-end.

Building My First Model	- Course overview - Introduction to Machine Learning - Introduction to BigML's platform - Building my first ML model - Summary and next session
Supervised Learning – Part I	- Key algorithms - Evaluations - Ensembles - Linear Regressions - Logistic Regressions - Real-world examples - Summary and next session
Supervised Learning – Part II	- Deepnets - OptiML - Fusions - Advanced Evaluations, overfitting, and leakage - Real-world examples - Summary and next session
Unsupervised Learning – Part I	- Time Series - Anomaly Detection - Cluster Analysis - Real-world examples combined with supervised learning algorithms - Summary and next session
Unsupervised Learning – Part II	- Association Discovery - Topic Modeling - Principal Component Analysis (PCA) - Real-world examples combined with supervised learning algorithms - Summary and next session
Data Preparation	- Data transformations - Feature engineering - Feature selection - Summary and next session
End-to-End Machine Learning	- ML workflows - Programming BigML - BigML Rest API and Bindings - BigML command line tool: BigMLer - Server-side scripts: WhizzML - Summary and next session

Image Recognition	- Structured vs unstructured data
	- Traditional techniques
	- Convolutional Neural Networks (CNN)
	- CNN architectures
	- Training Convolutional Neural Networks
	  - Image classification vs object detection
	  - Transfer learning techniques
	- Practice what you have learned