

A wonderful introduction to graphs for machine learning enthusiasts, as well as a great entrée into machine learning for graph experts.

Erik Sapper

At its core, machine learning is about efficiently identifying patterns and relationships in data. Many tasks, such as finding associations among terms so you can make accurate search recommendations or locating individuals within a social network who have similar interests, are naturally expressed as graphs. *Graph-Powered Machine Learning* teaches you how to use graph-based algorithms and data organization strategies to develop superior machine learning applications.

table of contents

brief table of contents

PART 1: INTRODUCTION

1 MACHINE LEARNING AND GRAPH: AN INTRODUCTION

- 1.1 Introduction to Machine Learning
 - 1.1.1 Machine Learning Project Lifecycle
 - 1.1.2 Algorithm taxonomies
- 1.2 Machine Learning Challenges
 - 1.2.1 The source of truth
 - 1.2.2 Performance
 - 1.2.3 Storing the model
 - 1.2.4 Real time
- 1.3 Graphs
 - 1.3.1 What is a graph?
 - 1.3.2 Representing Graphs
 - 1.3.3 Graph as model of networks
 - 1.3.4 Property graph model
- 1.4 The role of graph in the machine learning
 - 1.4.1 Data Management
 - 1.4.2 Data Analysis
 - 1.4.3 Data Visualization
- 1.5 Summary
- 1.6 References

2 GRAPH DATA ENGINEERING

- 2.1 Working with Big Data
 - 2.1.1 Volume
 - 2.1.2 Velocity
 - 2.1.3 Variety

- 2.1.4 Veracity
- 2.2 Graphs in Big Data Platform
 - 2.2.1 Graphs are valuable for big data
 - 2.2.2 Graph for master data management
- 2.3 Graph Databases
 - 2.3.1 Graph Data Management
 - 2.3.2 Sharding
 - 2.3.3 Replication
 - 2.3.4 Native graph vs not native graph
 - 2.3.5 Neo4j
- 2.4 Summary
- 2.5 References

3 GRAPHS IN MACHINE LEARNING APPLICATION

- 3.1 Graphs in Machine Learning Workflow
 - 3.1.1 Managing Data Sources
 - 3.1.2 Algorithms
 - 3.1.3 Storing and Accessing Machine Learning Model
 - 3.1.4 Visualization
- 3.2 Graph as processing pattern
 - 3.2.1 Pregel
- 3.3 Graph for defining complex processing workflow
 - 3.3.1 DataFlow
- 3.4 Summary
- 3.5 References

PART 2: TECHNIQUES

4 DATA MODELING

5 LINK ANALYSIS

6 GRAPH CLUSTERING

7 GRAPH CLASSIFICATION

8 NEAREST NEIGHBOR

PART 3: APPLICATIONS

9 NATURAL LANGUAGE PROCESSING

10 RECOMMENDATION

11 RELEVANT SEARCHING ON A KNOWLEDGE GRAPH

APPENDIXES

APPENDIX A: BIG DATA ARCHITECTURE

APPENDIX B: TOOLKIT

about the technology

Graph-based machine learning is an incredibly powerful tool for any task that involves pattern matching in large data sets. Applications include security concerns like identifying fraud or detecting network intrusions, application areas like social networking or natural language processing, and better user experiences through accurate recommendations and smart search. By organizing and analyzing your data as graphs your applications work more fluidly with graph-centric algorithms like nearest neighbor or page rank where it's important to quickly identify and exploit relevant relationships. Modern graph data stores, like Neo4j or Amazon Neptune, are readily available tools that support graph-powered machine learning.

about the book

Graph-Powered Machine Learning introduces you to graph technology concepts, highlighting the role of graphs in machine learning and big data platforms. You'll get an in-depth look at techniques including data source modeling, algorithm design, link analysis, classification, and clustering. As you master the core concepts, you'll explore three end-to-end projects that illustrate architectures, best design practices, optimization approaches, and common pitfalls. Author Alessandro Negro's extensive experience building graph-based machine learning systems shines through in every chapter, as you learn from examples and concrete scenarios based on his own work with real clients!

what's inside

- The lifecycle of a machine learning project
- Three end-to-end applications
- Graphs in big data platforms
- Data source modeling
- Natural language processing, recommendations, and relevant search
- Optimization methods

about the reader

Written for readers comfortable with machine learning basics.

about the author

Alessandro Negro is a Chief Scientist at GraphAware. With extensive experience in software development, software architecture, and data management, he has been a speaker at many conferences, such as Java One, Oracle Open World, and Graph Connect. He holds a Ph.D. in Computer Science and has authored several publications on graph-based machine learning.