Personalized Machine Learning

Every day we interact with machine learning systems offering individualized predictions for our entertainment, social connections, purchases, or health. These involve several modalities of data, from sequences of clicks to text, images, and social interactions. This book introduces common principles and methods that underpin the design of personalized predictive models for a variety of settings and modalities.

The book begins by revising 'traditional' machine learning models, focusing on how to adapt them to settings involving user data; then presents techniques based on advanced principles such as matrix factorization, deep learning, and generative modeling; and concludes with a detailed study of the consequences and risks of deploying personalized predictive systems.

A series of case studies in domains ranging from e-commerce to health plus hands-on projects and code examples will give readers understanding and experience with large-scale real-world datasets and the ability to design models and systems for a wide range of applications.

JULIAN MCAULEY has been a Professor at the University of California San Diego since 2014. Personalized Machine Learning is the main research area of his lab, with applications ranging from personalized recommendation to dialog, health care, and fashion design. He regularly collaborates with industry on these topics, including Amazon, Facebook, Microsoft, Salesforce, and Etsy. His work has been selected for several awards including an NSF CAREER award, and faculty awards from Amazon, Salesforce, Facebook, and Qualcomm, among others.

Personalized Machine Learning

JULIAN MCAULEY
University of California San Diego

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

103 Penang Road, #05–06/07, Visioncrest Commercial, Singapore 23846

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781316518908
DOI: 10.1017/9781009003971

First published 2022

A catalogue record for this publication is available from the British Library.

ISBN 978-1-316-51890-8 Hardback

Contents

Cambridge University Press
978-1-316-51890-8 — Personalized Machine Learning
Julian McAuley
Frontmatter
More Information

Notation

Common Mathematical Symbols

Machine Learning

y	vector of labels
X	matrix of features
x_i	feature vector for the ith sample
$f(x_i)$	model prediction for the ith sample
r_i	residual (error) associated with the ith prediction, $r_i = (y_i - f(x_i))$
θ	vector of model parameters
σ	sigmoid function $\sigma(x) = \frac{1}{1+e^{-x}}$
$\|\|x\|\|_p$	p-norm, $\|\|x\|\|_p = (\sum_i \|x_i\|^p)^{1/p}$
ℓ_1; ℓ_2	regularizers $\|\|\theta\|\|_1$ and $\|\|\theta\|\|_2$
λ	regularization hyperparameter
$\mathcal{L}$; ℓ	likelihood and log-likelihood

Users and Items

$u \in U$	user u in user set U
$i \in I$	item i in item set I
I_u	set of items rated (or interacted with) by user u
U_i	set of users who have rated (or interacted with) item i
$\|U\|$; $\|I\|$	number of users and number of items
$R_{u,i}$	measurement (e.g., a rating) associated with an interaction between user u and item i
$x_{u,i}$	model estimate of the compatibility between user u and item i

Recommender Systems

β_u	bias term associated with user u
β_i	bias term associated with item i

γ_u	vector of parameters describing a single user u
γ_i	vector of parameters describing a single item i
γ_U; γ_I	parameters for all users U or all items I
K	feature dimensionality (or number of latent factors)

Common Abbreviations

AUC	area under the ROC curve (eq. (5.26))
BER	balanced error rate (eq. (3.20))
BPR	Bayesian personalized ranking (sec. 5.2.2)
CNN	convolutional neural network (sec. 5.5.4)
FVU	fraction of variance unexplained (eq. (2.32))
FN/FNR	false negatives/false negative rate (sec. 3.3.1)
FP/FPR	false positives/false positive rate (sec. 3.3.1)
GAN	generative adversarial network (sec. 9.4)
LSTM	long short-term memory model (sec. 7.6)
MAE	mean absolute error (eq. (2.17))
MLE	maximum likelihood estimation (sec. 2.2.3)
MLP	multilayer perceptron (sec. 5.5.2)
MMR	maximal marginal relevance (sec. 10.3.1)
MRR	mean reciprocal rank (sec. 5.4.2)
MSE	mean squared error (sec. 2.2.1)
NDCG	normalized discounted cumulative gain (sec. 5.4.3)
RNN	recurrent neural network (sec. 7.6)
ROC	receiver-operating characteristic (sec. 3.3.3)
SVM	support vector machine (sec. 3.2)
TF-IDF	term frequency and inverse document frequency (eq. (8.8))
TN/TNR	true negatives/true negative rate (sec. 3.3.1)
TP/TPR	true positives/true positive rate (sec. 3.3.1)