

# Wiley Series on Parallel and Distributed Computing

Albert Y. Zomaya, Series Editor

Diane J. Cook • Narayanan C. Krishnan

# ACTIVITY LEARNING

*Discovering, Recognizing, and Predicting  
Human Behavior from Sensor Data*

WILEY

# *Activity Learning*

*Discovering, Recognizing,  
and Predicting Human Behavior  
from Sensor Data*

**Diane J. Cook  
Narayanan C. Krishnan**

**WILEY**

Copyright © 2015 by John Wiley & Sons, Inc. All rights reserved

Published by John Wiley & Sons, Inc., Hoboken, New Jersey

Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at [www.copyright.com](http://www.copyright.com). Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

**Limit of Liability/Disclaimer of Warranty:** While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at [www.wiley.com](http://www.wiley.com).

***Library of Congress Cataloging-in-Publication Data:***

Cook, Diane J., 1963-

Activity learning : discovering, recognizing, and predicting human behavior from sensor data / Diane J. Cook, Narayanan C. Krishnan.

pages cm. – (Wiley series on parallel and distributed computing)

Includes bibliographical references and index.

ISBN 978-1-118-89376-0 (hardback)

1. Active learning—Data processing. 2. Detectors—Data processing. 3. Multisensor data fusion. I. Title. LB1027.23.C665 2015  
371.3—dc23

2014039501

Typeset in 10/12pt TimesLTStd by Laserwords Private Limited, Chennai, India.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

1 2015

## *Activity Learning*

---

**WILEY SERIES ON PARALLEL  
AND DISTRIBUTED COMPUTING**

**Series Editor: Albert Y. Zomaya**

A complete list of titles in this series appears at the end of this volume.

---

# *Preface*

Activity learning from sensor data is a research topic that is making its way into many fields including machine learning, pervasive computing, psychology, and sociology. The ability to discover and model activities based on collected sensor data has matured. These methods are now able to handle increasingly complex with increasingly complex situations including unscripted activities, interrupted or interwoven activities, real-time learning, and learning from multiple users or residents. At the same time, there have been advances in the application of the theoretical techniques to challenge real-world problems including health monitoring and home automation.

The goal of this book is to define the notion of an activity model learned from sensor data and to present the key algorithms that form the core of the field. While many current efforts have contributed to the ability to automatically recognize known activities, there are other aspects of activity learning that we want to include in this treatment. These include discovering activity patterns from unlabeled data and predicting when specific activities will occur in the future, as well as mapping sensor event sequences to predefined activity labels. We also want to discuss the challenges that are faced when these theoretical techniques are applied to real-world problems and suggest methods for addressing the challenges.

This book is designed for an interdisciplinary audience who would like to use or design activity learning techniques. As a result, most of the ideas are presented from the ground up, with little assumptions about the background of the reader. Ideally, the book will provide some helpful background and guidance to researchers, undergraduate or graduate students, or practitioners who want to incorporate the ideas into their own work.

The best way to understand activity learning techniques is to look at activity sensor data, play with existing tools, and writing your own code. To help with this process, we provide code for many of the techniques described in this book. A variety of

activity sensor datasets are available for your use as well. All of the related book materials can be found online at <http://eecs.wsu.edu/~cook/albook>.

We would like to thank the many individuals who helped us with this book. Jacqueline Southwick took the activity photographs that we included and the team of Kyle Elsalhi and Anthony Simmons, who generated the Kinect-based motion history image. Numerous individuals contributed to the data collections we made available with this book and who provided feedback on drafts of the book. These include Larry Holder, Maureen Schmitter-Edgecombe, Aaron Crandall, Brian Thomas, Yasmin Sahaf, Kyle Feuz, Prafulla Dawadi, Ehsan Nazerfard, Bryan Minor, Adriana Seelye, Carolyn Parsey, Jennifer Walker, Alyssa Weakley, Sue Nelson, Thomas Cowger, Selina Akter, and Prasanth Lade. Our editors at Wiley provided guidance throughout the entire process and for that we are grateful. Last but not least, we want to thank our colleagues, friends, and family who were unfailingly supportive of this effort. Diane would like to thank her family, Larry, Abby, and Ryan, who kept her going with ideas, encouragement, and lots of humor. She dedicates this book to them. Narayanan dedicates this book to his family, Krishnan, Geetha, and Karthik for their constant encouragement and support.

DIANE J. COOK and  
NARAYANAN C. KRISHNAN

# *List of Figures*

Figure 1.1	The role of activity learning in the design of an intelligent agent	2
Figure 2.1	Examples of individual group actions and activities that are found in common everyday life settings	6
Figure 2.2	Relationship between state (S), action, and activity	7
Figure 3.1	PIR infrared sensor packaging (left) and internal architecture (right)	12
Figure 3.2	The two components of a magnetic contact switch encased in plastic (left) and a magnetic contact switch acting as a door open/shut sensor (right)	13
Figure 3.3	Tactile pressure sensors positioned in a chair (left) and interpreted as a pressure map (right)	14
Figure 3.4	RFID tag positioned beneath visible object barcode	15
Figure 3.5	Increased acceleration values can indicate a start or stop of a motion	16
Figure 3.6	A smart phone with a three-axis accelerometer (left) and a three-axis gyroscope (right)	16



Figure 3.7	Floor plan and sensor layout for the first floor of a smart home. Sensor identifiers starting with “M” are infrared motion sensors, sensors starting with “D” are magnetic door closure sensors, and sensors starting with “T” are ambient temperature sensors. The hand soap object sensor is located near the sink and the additional object sensors are in the kitchen closet. All object sensors are indicated with stars. Additional sensors included in these datasets are A001 (gas usage sensor attached to burner), A002 and A003 (hot and cold water consumption sensors, respectively), and P001 (whole-home power usage meter)	18
Figure 3.8	Positioning of wearable accelerometers on an individual’s dominant arm (left) and hip (right)	19
Figure 3.9	Time plot of sensor activity during the Hand Washing activity. The x axis shows the time of day when the sensor message was generated. The y axis shows the identifier (on the left) and functional location (on the right) for each sensor generating messages	19
Figure 3.10	Time plot of normalized accelerometer readings during the Hand Washing activity. The x axis shows the time of day when the sensor message was generated. The y axis shows the identifier (on the left) and placement (on the right) for each sensor. There are three lines for each sensor, corresponding to the x, y, and z axes	19
Figure 3.11	Time plot of sensor activity during the Sweeping activity. The x axis shows the time of day when the sensor message was generated. The y axis shows the identifier (on the left) and functional location (on the right) for each sensor generating messages	20
Figure 3.12	Time plot of normalized accelerometer readings during the Sweeping activity. The x axis shows the time of day when the sensor message was generated. The y axis shows the identifier (on the left) and placement (on the right) for each sensor. There are three lines for each sensor, corresponding to the x, y, and z axes	20
Figure 3.13	Heat map of discrete event sensor usage in a smart home (left, motion sensor locations represented by circles in the image) and motion history image of a sit-to-stand transition movement (right). In both images, the pixel intensities indicate activities occurring in the corresponding region of the space. In the left picture, darker circles indicate areas in which more time was spent. In the right picture, lighter regions indicate more recent movement occurring in the region	24

Figure 3.14	Plot of acceleration ( $A_x$ , $A_y$ , $A_z$ ) and rotational velocity ( $R_x$ , $R_y$ , $R_z$ ) values for HIP accelerometer from the Sweeping activity data sample	26
Figure 3.15	Frequency content in the accelerometer and gyroscope signal for the sensor placed on the hip while performing a Sweeping activity	33
Figure 4.1	Plot illustrating the distribution of the start time of the Eating activity from a synthetic dataset. The three solid curves represent the distribution for breakfast, lunch, and dinner. The dashed curve is the sum of these distributions	48
Figure 4.2	The GMM_EM algorithm	50
Figure 4.3	Graphical representation of an HMM. The shaded nodes represent the observed variables—the sensor events or the feature vectors from the sensor stream. The white nodes represent the hidden variables, which correspond to the underlying activity labels. Note the direction of the arrow indicating that the underlying hidden states (activities) generate the observations (sensor events)	51
Figure 4.4	A HMM used to recognize four distinct activities	51
Figure 4.5	The Viterbi algorithm to generate the most likely sequence of activity labels from a sequence of sensor observations	53
Figure 4.6	A sample decision tree for classifying a data point as Bed Toilet Transition, Personal Hygiene, Take Medicine, or Other	54
Figure 4.7	ID3 decision tree algorithm	55
Figure 4.8	Illustration of a SVM for a linearly separable classification problem	57
Figure 4.9	Illustration of the soft-margin concept for SVM classification in the nonseparable scenario	59
Figure 4.10	The two-class SVM algorithm	61
Figure 4.11	Graphical representation of a linear chain CRF	62
Figure 4.12	The Boosting algorithm to combine multiple classifiers	65
Figure 4.13	The Bagging algorithm to combine classifiers	66
Figure 4.14	The filter-based feature selection algorithm	68
Figure 4.15	The PCA dimensionality reduction algorithm	70
Figure 4.16	Plot of the data points projected onto a three-dimensional subspace using PCA. The axes in the figure represent the three principal components	72

Figure 5.1	The AR process includes the stages of raw sensor data collection, sensor preprocessing and segmentation, feature extraction and selection, classifier training, and data classification	76
Figure 5.2	Illustration of alternative approaches to processing sensor data streams for AR. Sensor event types and timings are depicted by the vertical lines, where the line type indicates the type, location, and value of the sensor message. The sensor windows are obtained using explicit segmentation, sliding window extraction with a fixed window time duration, or sliding window extraction with a fixed window length (number of sensor events)	77
Figure 5.3	Rule-based activity segmentation algorithm	80
Figure 5.4	Illustration of time dependency in a sliding window of sensor events	83
Figure 5.5	Effect of $X$ on weights	84
Figure 5.6	Illustration of sensor dependency in a sliding window of sensor events	85
Figure 5.7	Intensity-coded MI between every pair of motion sensors (right) in a smart home (left)	87
Figure 5.8	Online change point detection algorithm	91
Figure 5.9	Hybrid bottom-up/sliding-window segmentation algorithm	93
Figure 5.10	Example ROC curve (left) and PRC (right)	99
Figure 5.11	Event-based evaluation of Hand Washing AR. The vertical solid lines indicate the performance segment boundaries. The top row shows the ground truth labeling of hand washing or sweeping and the bottom row shows the labels generated by an AR algorithm. Each segment mismatch is labeled with the type of error that is represented, either an overfill (O), underfill (U), insertion (I), deletion (D), merge (M), or fragmenting (F)	100
Figure 6.1	Time spent by Americans as reported by the Bureau of Labor Statistics <sup>14</sup>	108
Figure 6.2	Example activity-attribute matrix. Each row represents a distinct activity and each column represents a separate high-level feature. The contents of the matrix indicate whether the feature is associated with the activity (value=1) or not (value=0)	109
Figure 6.3	Frequency-based sequence mining algorithm	112
Figure 6.4	Highest-value patterns found by frequency mining	113

Figure 6.5	Example of the hierarchical discovery algorithm applied to an input sequence $S$ . Each unique sensor event symbol is represented by a unique circle pattern in the diagram. A sequence pattern ( $P$ ) is identified and used to compress the dataset into a smaller sequence $S^1$ . A new best pattern ( $P^1$ ) is found in the second iteration of the hierarchical discovery algorithm	115
Figure 6.6	Hierarchical compression-based sequence mining algorithm	116
Figure 6.7	Illustration of the topic model process. Each sensor event sequence is a mixture of topics (high-level features), each topic is a distribution over low-level sensor features, each sensor feature is drawn from one of the topics	120
Figure 7.1	LZ78 string parsing algorithm	129
Figure 7.2	The trie formed by the LZ78 parsing of the location sequence aaababbbbbaabccddcbaaaa. The numbers in parentheses indicate the frequency of the corresponding phrase encountered within the sequence as it is parsed	130
Figure 7.3	ALZ string parsing algorithm	131
Figure 7.4	The trie formed by the ALZ parsing of the location sequence aaababbbbbaabccddcbaaaa	132
Figure 7.5	Forecasted time elapsed values for Enter Home activity	134
Figure 7.6	A regression tree to output expected number of seconds until the next occurrence of a Bed Toilet Transition activity. Each internal node queries the value of a particular variable and each node contains a multivariate linear regression function. Each of the regression functions generates a prediction value as a function of the lags and other input feature values. The features include lag variables, or values of the predicted variable at previous time units. Lag1 indicates the prediction value at time $t - 1$ and Lag3 indicates the prediction value at time $t - 3$	136
Figure 7.7	An order 2 AR-HMM	138
Figure 7.8	Actual and predicted class values over a three day period based on a regression tree model. In this graph, the x-axis represents the time of day (from the beginning of day one to the end of day three) and the y-axis represents the number of seconds (predicted or actual) that will elapse before the next occurrence of the target activity, Bed Toilet Transition. The class variable is the number of seconds that will elapse before the next occurrence of the target activity	147

Figure 8.1	The VARS video annotation tool allows a researcher to label aspects of an activity scene while watching the corresponding video	150
Figure 8.2	An experimenter observes a participant performing everyday activities via web cameras and provides activity labels remotely using the RAT	151
Figure 8.3	Example tweets that indicate an individual's activities	152
Figure 8.4	Illustration of a social network used to train an activity classifier for user $u_1$	153
Figure 8.5	Example of SMOTE synthetic data point creation for $k = 5$ nearest neighbors of point $x_i$	155
Figure 8.6	GAL generalized active learning algorithm	157
Figure 8.7	Learning curves for two algorithms, A and B. Algorithm A converges to optimal performance more quickly and thus has a larger ALC value than Algorithm B	158
Figure 8.8	Depiction of difference between supervised learning and transfer learning approaches to activity modeling	159
Figure 8.9	The TrAdaBoost algorithm to transfer instances from the source to target domains	162
Figure 8.10	Sample WikiHow pages for the activities Cook, Shower, and Jog	165
Figure 8.11	The CDAR algorithm to transfer instances between activity labels	166
Figure 8.12	Multi-view transfer learning	168
Figure 8.13	The Co-Training algorithm	168
Figure 8.14	The Co-EM algorithm	169
Figure 8.15	The Teacher-Learner algorithm	170
Figure 8.16	The Manifold Alignment algorithm	171
Figure 8.17	The number of distinct activity labels associated with data points sampled from a two-resident smart home	172
Figure 8.18	The binary relevance multi-label learning algorithm	174
Figure 8.19	The principal label space transform multi-label learning algorithm	176
Figure 8.20	The multi-label output code algorithm using canonical correlation analysis	179
Figure 8.21	Graphical representation of CHMM model for two users, denoted $a$ and $b$	185

Figure 8.22	Factored CRF representing a sequence of activities for two users, $a$ and $b$	186
Figure 8.23	Flow network for three sensor events with one cluster (path) highlighted	188
Figure 8.24	The STAR particle filter-based algorithm for tracking multiple users	191
Figure 9.1	The COACH hand-washing assistance system (a) Activity forecasting-based prompts can be delivered to a computer, television (b), or mobile device (c). Source: Reproduced by permission of Alex Mihailidis	197
Figure 9.2	A “Them and Us” game replayed inside the after-game visualization interface. Source: Reproduced by permission of Alan Chamberlin.	199
Figure 9.3	Example hierarchical model learned from MavHome sensor data	200
Figure 9.4	Activity profiles based on visualization of accelerometer data. Source: Copyright©Elsevier, 2015	202
Figure 9.5	Sample activity density maps	203
Figure 9.6	Radial bar chart for the Bed Toilet Transition activity	203
Figure 9.7	Visualization of activity density, activity times, and activity trends	204
Figure 9.8	Visualization of an activity daily diary	205
Figure 9.9	Wi-Fi cluster co-occurrence plot. Source: Copyright© Elsevier, 2015	206
Figure 9.10	Visualization of how surveyed Americans spent their time in 2008. Source: Reproduced by permission of IOP Publishing	207
Figure 9.11	Mobility patterns grouped by cell tower region for two time periods. The bars on the right correspond to the number of calls per hour on a log scale	208
Figure 9.12	Plot of relative activity level as a function of the hour of the day (left) and relative activity duration as a function of the activity class (right)	209
Figure 9.13	The meaningfulness of a place is indicated by friendships that are made there. Source: Reproduced by permission of Alex Pentland	210

# *Contents*

<b>Preface</b>	<b>ix</b>
<b>List of Figures</b>	<b>xi</b>
<b>1. Introduction</b>	<b>1</b>
<b>2. Activities</b>	<b>5</b>
2.1 Definitions / 5	
2.2 Classes of Activities / 7	
2.3 Additional Reading / 8	
<b>3. Sensing</b>	<b>11</b>
3.1 Sensors Used for Activity Learning / 11	
3.1.1 Sensors in the Environment / 12	
3.1.2 Sensors on the Body / 15	
3.2 Sample Sensor Datasets / 17	
3.3 Features / 17	
3.3.1 Sequence Features / 21	
3.3.2 Discrete Event Features / 23	
3.3.3 Statistical Features / 25	
3.3.4 Spectral Features / 31	
3.3.5 Activity Context Features / 34	

3.4	Multisensor Fusion / 34	
3.5	Additional Reading / 38	
<b>4.</b>	<b>Machine Learning</b>	<b>41</b>
4.1	Supervised Learning Framework / 41	
4.2	Naïve Bayes Classifier / 44	
4.3	Gaussian Mixture Model / 48	
4.4	Hidden Markov Model / 50	
4.5	Decision Tree / 54	
4.6	Support Vector Machine / 56	
4.7	Conditional Random Field / 62	
4.8	Combining Classifier Models / 63	
4.8.1	Boosting / 64	
4.8.2	Bagging / 65	
4.9	Dimensionality Reduction / 66	
4.10	Additional Reading / 72	
<b>5.</b>	<b>Activity Recognition</b>	<b>75</b>
5.1	Activity Segmentation / 76	
5.2	Sliding Windows / 81	
5.2.1	Time Based Windowing / 81	
5.2.2	Size Based Windowing / 82	
5.2.3	Weighting Events Within a Window / 83	
5.2.4	Dynamic Window Sizes / 87	
5.3	Unsupervised Segmentation / 88	
5.4	Measuring Performance / 92	
5.4.1	Classifier-Based Activity Recognition Performance Metrics / 95	
5.4.2	Event-Based Activity Recognition Performance Metrics / 99	
5.4.3	Experimental Frameworks for Evaluating Activity Recognition / 102	
5.5	Additional Reading / 103	
<b>6.</b>	<b>Activity Discovery</b>	<b>107</b>
6.1	Zero-Shot Learning / 108	



6.2	Sequence Mining / 110	
6.2.1	Frequency-Based Sequence Mining / 111	
6.2.2	Compression-Based Sequence Mining / 112	
6.3	Clustering / 117	
6.4	Topic Models / 119	
6.5	Measuring Performance / 121	
6.5.1	Expert Evaluation / 121	
6.6	Additional Reading / 124	
<b>7.</b>	<b>Activity Prediction</b>	<b>127</b>
7.1	Activity Sequence Prediction / 128	
7.2	Activity Forecasting / 133	
7.3	Probabilistic Graph-Based Activity Prediction / 137	
7.4	Rule-Based Activity Timing Prediction / 139	
7.5	Measuring Performance / 142	
7.6	Additional Reading / 146	
<b>8.</b>	<b>Activity Learning in the Wild</b>	<b>149</b>
8.1	Collecting Annotated Sensor Data / 149	
8.2	Transfer Learning / 158	
8.2.1	Instance and Label Transfer / 162	
8.2.2	Feature Transfer with No Co-occurrence Data / 166	
8.2.3	Informed Feature Transfer with Co-occurrence Data / 167	
8.2.4	Uninformed Feature Transfer with Co-occurrence Data Using a Teacher-Learner Model / 168	
8.2.5	Uninformed Feature Transfer with Co-occurrence Data Using Feature Space Alignment / 170	
8.3	Multi-Label Learning / 170	
8.3.1	Problem Transformation / 173	
8.3.2	Label Dependency Exploitation / 174	
8.3.3	Evaluating the Performance of Multi-Label Learning Algorithms / 179	
8.4	Activity Learning for Multiple Individuals / 180	
8.4.1	Learning Group Activities / 180	
8.4.2	Train on One/Test on Multiple / 183	