


Applied Affective Computing

Leimin Tian
Sharon Oviatt
Michal Muszynski
Brent C. Chamberlain
Jennifer Healey
Akane Sano



ASSOCIATION FOR COMPUTING MACHINERY

Applied Affective Computing

ACM Books

Editors in Chief

Sanjiva Prasad, *Indian Institute of Technology (IIT) Delhi, India*

Marta Kwiatkowska, *University of Oxford, UK*

Charu Aggarwal, *IBM Corporation, USA*

ACM Books is a new series of high-quality books for the computer science community, published by ACM in collaboration with Morgan & Claypool Publishers. ACM Books publications are widely distributed in both print and digital formats through booksellers and to libraries (and library consortia) and individual ACM members via the ACM Digital Library platform.

Circuits, Packets, and Protocols: Entrepreneurs and Computer Communications, 1968–1988

James L. Pelkey

Andrew L. Russell, *SUNY Polytechnic Institute, New York*

Loring G. Robbins

2022

Theories of Programming: The Life and Works of Tony Hoare

Editors: Cliff B. Jones, *Newcastle University, UK*

Jayadev Misra, *The University of Texas at Austin, US*

2021

Software: A Technical History

Kim W. Tracy, *Rose-Hulman Institute of Technology, IN, USA*

2021

The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition

Editors: Birgit Lugin, *Julius-Maximilians-Universität of Würzburg*

Catherine Pelachaud, *CNRS-ISIR, Sorbonne Université*

David Traum, *University of Southern California*

2021

Probabilistic and Causal Inference: The Works of Judea Pearl

Editors: Hector Geffner, *ICREA and Universitat Pompeu Fabra*

Rina Dechter, *University of California, Irvine*

Joseph Y. Halpern, *Cornell University*

2022

Event Mining for Explanatory Modeling

Laleh Jalali, *University of California, Irvine (UCI), Hitachi America Ltd.*

Ramesh Jain, *University of California, Irvine (UCI)*

2021

Intelligent Computing for Interactive System Design: Statistics, Digital Signal Processing, and Machine Learning in Practice

Editors: Parisa Eslambolchilar, *Cardiff University, Wales, UK*

Andreas Komninos, *University of Patras, Greece*

Mark Dunlop, *Strathclyde University, Scotland, UK*

2021

Semantic Web for the Working Ontologist: Effective Modeling for Linked Data, RDFS, and OWL, Third Edition

Dean Allemang, *Working Ontologist LLC*

Jim Hendler, *Rensselaer Polytechnic Institute*

Fabien Gandon, *INRIA*

2020

Code Nation: Personal Computing and the Learn to Program Movement in America

Michael J. Halvorson, *Pacific Lutheran University*

2020

Computing and the National Science Foundation, 1950–2016: Building a Foundation for Modern Computing

Peter A. Freeman, *Georgia Institute of Technology*

W. Richards Adrion, *University of Massachusetts Amherst*

William Aspray, *University of Colorado Boulder*

2019

Providing Sound Foundations for Cryptography: On the work of Shafi Goldwasser and Silvio Micali

Oded Goldreich, *Weizmann Institute of Science*

2019

Concurrency: The Works of Leslie Lamport

Dahlia Malkhi, *VMware Research and Calibra*

2019

The Essentials of Modern Software Engineering: Free the Practices from the Method Prisons!

Ivar Jacobson, *Ivar Jacobson International*

Harold “Bud” Lawson, *Lawson Konsult AB (deceased)*

Pan-Wei Ng, *DBS Singapore*

Paul E. McMahon, *PEM Systems*

Michael Goedicke, *Universität Duisburg–Essen*

2019

Data Cleaning

Ihab F. Ilyas, *University of Waterloo*

Xu Chu, *Georgia Institute of Technology*

2019

Conversational UX Design: A Practitioner's Guide to the Natural Conversation Framework

Robert J. Moore, *IBM Research–Almaden*

Raphael Arar, *IBM Research–Almaden*

2019

Heterogeneous Computing: Hardware and Software Perspectives

Mohamed Zahran, *New York University*

2019

Hardness of Approximation Between P and NP

Aviad Rubinfeld, *Stanford University*

2019

The Handbook of Multimodal-Multisensor Interfaces, Volume 3: Language Processing, Software, Commercialization, and Emerging Directions

Editors: Sharon Oviatt, *Monash University*

Björn Schuller, *Imperial College London and University of Augsburg*

Philip R. Cohen, *Monash University*

Daniel Sonntag, *German Research Center for Artificial Intelligence (DFKI)*

Gerasimos Potamianos, *University of Thessaly*

Antonio Krüger, *Saarland University and German Research Center for Artificial Intelligence (DFKI)*

2019

Making Databases Work: The Pragmatic Wisdom of Michael Stonebraker

Editor: Michael L. Brodie, *Massachusetts Institute of Technology*

2018

The Handbook of Multimodal-Multisensor Interfaces, Volume 2: Signal Processing, Architectures, and Detection of Emotion and Cognition

Editors: Sharon Oviatt, *Monash University*

Björn Schuller, *University of Augsburg and Imperial College London*

Philip R. Cohen, *Monash University*

Daniel Sonntag, *German Research Center for Artificial Intelligence (DFKI)*

Gerasimos Potamianos, *University of Thessaly*

Antonio Krüger, *Saarland University and German Research Center for Artificial Intelligence (DFKI)*

2018

Declarative Logic Programming: Theory, Systems, and Applications

Editors: Michael Kifer, *Stony Brook University*

Yanhong Annie Liu, *Stony Brook University*

2018

The Sparse Fourier Transform: Theory and Practice

Haitham Hassanieh, *University of Illinois at Urbana-Champaign*

2018

The Continuing Arms Race: Code-Reuse Attacks and Defenses

Editors: Per Larsen, *Immunant, Inc.*

Ahmad-Reza Sadeghi, *Technische Universität Darmstadt*

2018

Frontiers of Multimedia Research

Editor: Shih-Fu Chang, *Columbia University*

2018

Shared-Memory Parallelism Can Be Simple, Fast, and Scalable

Julian Shun, *University of California, Berkeley*

2017

Computational Prediction of Protein Complexes from Protein Interaction Networks

Sriganesh Srihari, *The University of Queensland Institute for Molecular Bioscience*

Chern Han Yong, *Duke-National University of Singapore Medical School*

Limsoon Wong, *National University of Singapore*

2017

The Handbook of Multimodal-Multisensor Interfaces, Volume 1: Foundations, User Modeling, and Common Modality Combinations

Editors: Sharon Oviatt, *Incaa Designs*

Björn Schuller, *University of Passau and Imperial College London*

Philip R. Cohen, *Voicebox Technologies*

Daniel Sonntag, *German Research Center for Artificial Intelligence (DFKI)*

Gerasimos Potamianos, *University of Thessaly*

Antonio Krüger, *Saarland University and German Research Center for Artificial Intelligence (DFKI)*

2017

Communities of Computing: Computer Science and Society in the ACM

Thomas J. Misa, Editor, *University of Minnesota*

2017

Text Data Management and Analysis: A Practical Introduction to Information Retrieval and Text Mining

ChengXiang Zhai, *University of Illinois at Urbana-Champaign*

Sean Massung, *University of Illinois at Urbana-Champaign*

2016

An Architecture for Fast and General Data Processing on Large Clusters

Matei Zaharia, *Stanford University*

2016

Reactive Internet Programming: State Chart XML in Action

Franck Barbier, *University of Pau, France*

2016

Verified Functional Programming in Agda

Aaron Stump, *The University of Iowa*

2016

The VR Book: Human-Centered Design for Virtual Reality

Jason Jerald, *NextGen Interactions*

2016

Ada's Legacy: Cultures of Computing from the Victorian to the Digital Age

Robin Hammerman, *Stevens Institute of Technology*

Andrew L. Russell, *Stevens Institute of Technology*

2016

Edmund Berkeley and the Social Responsibility of Computer Professionals

Bernadette Longo, *New Jersey Institute of Technology*

2015

Candidate Multilinear Maps

Sanjam Garg, *University of California, Berkeley*

2015

Smarter Than Their Machines: Oral Histories of Pioneers in Interactive Computing

John Cullinane, *Northeastern University; Mossavar-Rahmani Center for Business and Government, John F. Kennedy School of Government, Harvard University*

2015

A Framework for Scientific Discovery through Video Games

Seth Cooper, *University of Washington*

2014

Trust Extension as a Mechanism for Secure Code Execution on Commodity Computers

Bryan Jeffrey Parno, *Microsoft Research*

2014

Embracing Interference in Wireless Systems

Shyamnath Gollakota, *University of Washington*

2014

Applied Affective Computing

Leimin Tian

Monash University

Sharon Oviatt

Monash University

Michal Muszynski

Carnegie Mellon University and University of Geneva

Brent C. Chamberlain

Utah State University

Jennifer Healey

Adobe Research, San Jose

Akane Sano

Rice University

ACM Books #41



Copyright © 2022 by Association for Computing Machinery

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews—without the prior permission of the publisher.

Designations used by companies to distinguish their products are often claimed as trademarks or registered trademarks. In all instances in which the Association of Computing Machinery is aware of a claim, the product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

Applied Affective Computing

Leimin Tian, Sharon Oviatt, Michal Muszynski, Brent C. Chamberlain, Jennifer Healey and Akane Sano

books.acm.org

<http://books.acm.org>

ISBN: 978-1-4503-9590-8 hardcover

ISBN: 978-1-4503-9591-5 paperback

ISBN: 978-1-4503-9592-2 EPUB

ISBN: 978-1-4503-9593-9 eBook

Series ISSN: 2374-6769 print 2374-6777 electronic

DOIs:

10.1145/3502398 Book	10.1145/3502398.3502406 Chapter 7
10.1145/3502398.3502399 Preface	10.1145/3502398.3502407 Chapter 8
10.1145/3502398.3502400 Chapter 1	10.1145/3502398.3502408 Chapter 9
10.1145/3502398.3502401 Chapter 2	10.1145/3502398.3502409 Chapter 10
10.1145/3502398.3502402 Chapter 3	10.1145/3502398.3502410 Chapter 11
10.1145/3502398.3502403 Chapter 4	10.1145/3502398.3502411 Chapter 12
10.1145/3502398.3502404 Chapter 5	10.1145/3502398.3502412 Bibliography
10.1145/3502398.3502405 Chapter 6	10.1145/3502398.3502413 Bios/Index

A publication in the ACM Books series, #41

Editors in Chief: Sanjiva Prasad, *Indian Institute of Technology (IIT) Delhi, India*

Marta Kwiatkowska, *University of Oxford, UK*

Charu Aggarwal, *IBM Corporation, USA*

Area Editor: Michel Beaudouin-Lafon, *Université Paris-Sud*

This book was typeset in Arnhem Pro 10/14 and Flama using pdf_T_EX.

First Edition

10 9 8 7 6 5 4 3 2 1

In memory of Professor Patrizia Lombardo (1950–2019)

Contents

List of Figures xvii

List of Tables xxiii

Preface xxv

Acknowledgments xxvi

Chapter 1 Introduction to Applied Affective Computing 1

- 1.1 Affective Computing 1
- 1.2 Applied Affective Computing 4
- 1.3 Book Overview and Highlights 9
- 1.4 Contributions 13

Chapter 2 Emotions as Studied in Psychology and Cognitive Science 17

- 2.1 Theories of Emotion 17
- 2.2 Emotion as an Adaptive Function 22
- 2.3 Emotion as a Social Communicative Function 23
- 2.4 Summary 27

Chapter 3 Machine Learning Approaches for Applied Affective Computing 29

- 3.1 Machine Learning for Affective Computing 30
- 3.2 Deep Machine Learning for Affective Computing 31
- 3.3 Multimodal Representation of Affect: Knowledge-inspired versus Data-driven Features 32
- 3.4 Unimodal Features for Affect Recognition 32
- 3.5 Multimodal Emotion-aware Systems 40
- 3.6 Evaluation of Emotion-aware Systems 43
- 3.7 Synthesis of Emotion and Emotional Behaviors 44
- 3.8 Discussion 45
- 3.9 Conclusion 46

Chapter 4 Multimodal Data Collection and Processing for Applied Affective Computing 49

- 4.1 Multimodal Data Collection 49
- 4.2 Multimodal Data Processing 58
- 4.3 Multimodal Data Fusion 59
- 4.4 Conclusion and Future Work 65

Chapter 5 Emotion Recognition in the Wild 67

- 5.1 Modalities 68
- 5.2 Deep Learning for Emotion Recognition in the Wild 70
- 5.3 Emotion Representation in the Wild 72
- 5.4 Evaluation of Affect Recognition in the Wild 73
- 5.5 The Role of Context in Affect Recognition in the Wild 75
- 5.6 The Role of Environment in Affect Recognition in the Wild 75
- 5.7 Discussion 76
- 5.8 Conclusion 79

Chapter 6 Reinforcement Learning and Affective Computing 81

- 6.1 Reinforcement Learning 82
- 6.2 Affective Computing 88
- 6.3 Extending the Reinforcement Learning Framework 90
- 6.4 Conclusions 96

Chapter 7 Synthesizing Natural and Believable Emotional Expressions 99

- 7.1 Introduction 99
- 7.2 Emotional Expression Models 101
- 7.3 Synthesizing Emotional Expressions 102
- 7.4 The Role of Agency 103
- 7.5 Additional Considerations and Challenges 105

Chapter 8 Emotion-aware Human–Robot Interaction and Social Robots 107

- 8.1 Developing Emotion-aware Human–Robot Interaction Systems 107
- 8.2 Improving Social Intelligence of a Robot by Addressing Social Errors in Human–Robot Interaction 109
- 8.3 Case Study: Prompting Human Assistance in Human–Robot Collaboration Using Robots’ Emotional Expressions 110
- 8.4 Case Study: Understanding Users’ Expectations of Robots in Public Spaces 121
- 8.5 Challenges in Affective HRI 126
- 8.6 Guidelines for Developing Affective HRI Systems and Social Robots 129

Chapter 9 Affective Computing for Enhancing Well-Being 131

- 9.1 Motivation 131
- 9.2 Sensing and Analytics 133
- 9.3 Examples of Well-Being Studies 135
- 9.4 Sensing Beyond the Phone 139
- 9.5 Actions and Interventions 144
- 9.6 Conclusions 148

Chapter 10 Applied Affective Computing in Built Environments 153

- 10.1 Setting the Foundation Toward Emotionally Aware Planning and Design 153
- 10.2 Case Study: Passive Responses to Urban Infrastructure 156
- 10.3 Experimental Task 157
- 10.4 Data 159
- 10.5 Approach 161
- 10.6 Results 164
- 10.7 Conclusions 166
- 10.8 Guidelines for Affective Computing in Built Environments 167

Chapter 11 Addressing Ethical Issues of Affective Computing 169

- 11.1 Ethical Concerns of Affective Computing 169
- 11.2 A Fair System 172
- 11.3 A Privacy-preserving System 175
- 11.4 A Transparent System 178
- 11.5 A Beneficial System 180
- 11.6 A Responsible System 183
- 11.7 Conclusion 185

Chapter 12 Future of Affective Computing and Applied Affective Computing 187

- 12.1 Applied Affective Computing: Guidelines for Best Practice 187
- 12.2 Example Uses of the Guidelines 190
- 12.3 Open Challenges and Future Directions 194
- 12.4 Summary 203

Bibliography 205

Authors' Biographies 267

Index 271

List of Figures

1.1	An affective computing system.	3
1.2	Publication data on affective computing according to Web of Science.	4
1.3	A mindmap of applied affective computing.	9
2.1	Overlaps between the Big-6 model [Based on Ekman et al. 1972], the cross-cultural expressions model [Based on Cordaro et al. 2018], and the affective neuroscience model [Based on Panksepp and Watt 2011].	18
2.2	The circumplex model of emotion. Based on Russell [1980] .	20
2.3	The constructionism approach to emotions. Based on Lindquist [2013] .	21
3.1	The scheme of an emotion-aware system.	30
3.2	Natural emotional expressions to a fearful (on the left) and disgusting (on the right) video. The snapshots of the stimulus videos taken from the MAHNOB-HCI database show eye gaze overlaid and without eye gaze overlaid, frontal captured video, raw physiological signals, and raw eye gaze data. In the first row, the red circles show the fixation points and their radius indicates the time spent in each fixation point. The red lines indicate the moments where each of the snapshots was captured. Adapted with permission from Soleymani et al. [2011] .	34
3.3	Hierarchical fusion of audio-video movie features with physiological and behavioral reactions of spectators.	42

- 4.1 (a) A graphical representation of the circumplex model of affect with the horizontal axis representing the valence dimension and the vertical axis representing the arousal or activation dimension [Based on [Russell 1980](#)]. (b) Affective Slider, a digital self-reporting tool composed of two sliders that measure arousal (top) and pleasure (bottom) on a continuous scale. Adapted from [Betella and Verschure \[2016\]](#). Copyright: © 2016 Betella, Verschure. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. 51
- 4.2 Illustration of multimodal data fusion schemes (A) early fusion and (B) late fusion (SC: skin conductance, ST: skin temperature, AC: acceleration). Adapted with permission from [Li and Sano \[2020b\]](#). 60
- 4.3 The convolutional recurrent Multiple Kernel Learning model combining sentiment features in audio, video, and text. Adapted with permission from [Poria et al. \[2017\]](#). 63
- 4.4 (A) Schematics of locally connected Long Short-Term Memory Denoising Auto Encoder (LC-LSTM-DAE), a deep representation learning framework that learns and extracts physiological and behavioral features from raw wearable sensor data. (B) Correlation matrix between auto-learned features (AF-) and handcrafted features (CF-) pairs. The dimensionality corresponding to each modality/type of features were annotated in parenthesis. Annotated on the color bar is the adjusted significance threshold of ± 0.05 ($p < 2.3 \times 10^{-5}$). Adapted with permission from [Li and Sano \[2020a\]](#). 64
- 5.1 Selected frames from the Aff-Wild database showing different emotional states of humans: head poses and facial expressions with occlusions in different illumination conditions. Adapted with permission from [Kollias et al. \[2019\]](#). 69
- 5.2 The architecture of a CNN-based recognition model for arousal and valence prediction. Adapted with permission from [Kollias et al. \[2019\]](#). 71
- 5.3 The architecture of a multimodal and multitask deep machine model. Adapted with permission from [Kollias and Zafeiriou \[2019\]](#). 72

- 5.4 Arousal and valence space [Russell 1980]. Adapted with permission from Nogueira et al. [2015]. 73
- 5.5 The Aff-Wild2 database [Based on Kollias and Zafeiriou 2019]. (a) Unbalanced facial expression classes. (b) Unbalanced action unit (AU) classes. 74
- 5.6 Arousal and valence annotations of an Aff-Wild database video at the frame level contaminated with occlusions. Adapted with permission from Kollias et al. [2019]. 77
- 5.7 The graphical user interface of the annotation tool used to collect arousal/valence annotations of videos in the wild from the Aff-Wild database. Adapted with permission from Kollias et al. [2019]. 78
- 6.1 The reinforcement learning problem involves an agent takes an action a_t and receives both a reward for taking that action r_{t+1} (similar to the bandit problem) and a change in state (s_{t+1}) from the environment. The environment is outside of the agent's sphere of control. 83
- 6.2 In the multiarmed bandit problem a player tries to maximize the value of the cumulative reward it receives from successive "pulls" of the machine's "arm." Each pull draws a sample from the arm's probability distribution. Determining the "best" arm is easier when the distributions are as well separated as on the left. In the two-armed bandit on the right the best arm is difficult to identify; however, the choice matters less as the expected values are more similar. 84
- 6.3 In the full reinforcement learning problem the agent takes an action (e.g., move left or move right) and both receives an instantaneous reward and changes its state with respect to the environment. 85
- 6.4 The value of a state requires knowing the value of possible future actions from that state. Uncertainties in the agent's next action can come from randomness in the behavior policy, uncertainty in the agent's perception of the environment, or variability in the environment's reward response. 87
- 6.5 This figure illustrates how a discount factor of $\gamma = 0.1$ would change the agent's view of the environment shown in Figure 6.4. The effect of the discount factor is to narrow the agent's sphere of concern. 87

6.6	The actor–critic architecture based on Sutton and Barto's [2018] diagram (left), and an extension of the agent and environment that includes a separately modeled internal environment (right) based on Singh et al's [2004] architecture. In the actor–critic architecture, the colored outline represents the boundary of the agent. (Left) Author Created. This is a figure representation of the discussion in Section 13.5 of Sutton and Barto [2018] . (Right) Based on Singh et al. [2004] .	91
6.7	A diagram of a tiered affective model where the internal environment of the organism is envisioned as having separate layers for sensing, reflexes, feelings, basic emotions, and complex emotions. The agent resides in the cognition layer and makes a decision based on signals from each of the affective evaluative layers. The layers are envisioned to develop sequentially from the right to the left and the agent's degree of control is assumed to diminish from left to right.	92
7.1	Emotion framework. Adapted with permission from Grandjean et al. [2008] .	103
8.1	The human–robot failure taxonomy proposed by Honig and Oron-Gilad [2018] . Copyright © 2018 Honig and Oron-Gilad.	109
8.2	The Cozmo robot.	112
8.3	The appraisal-based emotional model of the robot. Adapted with permission from Zhou and Tian [2020] .	113
8.4	Mapping between emotion categories and the arousal–valence emotion space. Adapted with permission from Zhou and Tian [2020] .	114
8.5	Emotion transition module. Adapted with permission from Zhou and Tian [2020] .	115
8.6	Procedure of the tower construction game. Adapted with permission from Zhou and Tian [2020] .	116
8.7	Stage participants reported that they understood the robots' intention during the tower construction game.	118
8.8	Perception of Cozmo's expressions.	119
8.9	Examples of robots in public spaces.	122
8.10	An example of multirobot systems in public spaces: the Niska robotic ice cream bar in Melbourne.	123

- 8.11 Participatory design using the RIZE end-user programming interface (left) and robot simulator (right). Adapted with permission from [Tian and Oviatt \[2021\]](#). 125
- 8.12 Incorporating physical and spatial context by evaluating designs on physical robot. Adapted with permission from [Tian et al. \[2020\]](#). 125
- 9.1 Enhancing well-being is one of the loftiest goals for affective computing. Well-being involves the totality of a person's life, so algorithms that wish to understand people and motivate them toward greater health and happiness need to consider a wide variety of both contexts and time periods. Photo by EVERST ([stock.adobe.com](#)). 132
- 9.2 The majority of well-being-related studies use the smartphone as the primary sensing platform given that it is the one device users are most likely to carry throughout their day. Photo by contrast-werkstatt ([stock.adobe.com](#)). 133
- 9.3 Stationary compute offers an opportunity to analyze facial expressions, keystroke patterns, and programs in focus. Photo by Syda Productions ([stock.adobe.com](#)). 140
- 9.4 People frequently put their phones down while at home to engage in other activities. The home environment also offers an opportunity to observe people in a different social context. Photo by Vasyl ([stock.adobe.com](#)). 141
- 9.5 The car is an excellent platform for observing people both longitudinally along repeated daily drives as well as in response to rare extreme events. Photo by Daxiao Productions ([stock.adobe.com](#)). 142
- 9.6 Positive computing is a class of interventions that focus on enhancing eight pro-social positive emotions: love, gratitude, hope, joy, forgiveness, compassion, faith, and awe. Photo by Oleg Breslavtsev ([stock.adobe.com](#)). 148
- 10.1 Examples images from the study area. 158
- 10.2 All 12 zones. Route started at start/end and continued in a counterclockwise direction through each zone, returning to the end zone. Actual start location was in a hotel slightly south of start/end cutoff. Adapted with permission from [Whitaker \[2018\]](#). 160
- 10.3 Frequency of grass spaces within each zone. Adapted with permission from [Whitaker \[2018\]](#). 161

- 10.4 Frequency of public light infrastructure within each zone. Adapted with permission from [Whitaker \[2018\]](#). 162
- 11.1 Six stages to align AI with human values. This is a figure representation created by the author to capture textual description in Section 3 of the cited paper based on [[Gabriel 2020](#)]. 181
- 12.1 The lifecycle of an affective computing study or system development. 195

List of Tables

1.1	Human versus machine's accuracy in the emotion classification experiments of Beard et al. [2018]	2
4.1	Most interpreted auto features for skin conductance, skin temperature, and acceleration, as well as the top 3 highest correlated crafted features with associated correlation coefficients ($p < 2.3 \times 10^{-5}$)	64
8.1	Taxonomy of social errors in HRI	111
8.2	Artificial emotions and outcomes of human–multirobot collaboration	117
8.3	Perceived intensity and naturalness of expressions	119
10.1	Machine learning models explored for this case study	164
10.2	Model D using 3FCV as the explanatory model. Not all elements existed in a full range from highest to lowest number (e.g., Num Trees 2–10 has some gaps because these conditions did not exist)	165
11.1	Human rights considered in national AI strategies and recommendations to policymakers	172
12.1	Guidelines for applied affective computing	188

Preface

Affective computing is a research field focusing on understanding human emotions and incorporating such knowledge in artificial intelligent agents. It is a rapidly evolving field since the concept was proposed more than two decades ago. Great advancements have been achieved in affective computing, for example, automatic emotion recognition has reached human-level performance in some tasks. With the maturing of affective computing technologies, today the affective computing community is facing a major transformation, with research outcomes leaving the controlled lab environments and being adopted in real-world applications. This book attempts to capture this transformation by proposing a novel concept called *applied affective computing*.

We define applied affective computing as “research and application efforts that focus on understanding the effectiveness, outcomes, and personal and societal impact of affective computing technologies in real-life usage.” This brings attention to the challenges of applying research findings to real-life scenarios. We first offer readers an overview to the state of the art and emerging themes in the field using comprehensive review of existing approaches to affective computing systems and social signal processing. We then provide in-depth case studies of applied affective computing in various domains, such as social robotics and mental well-being, to help readers emplace the discussed affective computing technologies in contexts. Furthermore, we discuss ethical concerns related to affective computing and how to prevent misuse of the technology to encourage responsible research and application. Finally, we identify future directions of the field and summarize a set of guidelines for developing next-generation affective computing systems that are effective, safe, and human-centered.

To researchers and practitioners new to affective computing, we hope that this book will serve as an introduction to the field to support them in identifying new research topics or developing novel applications. To more experienced researchers and practitioners, we hope that the discussions in this book will guide them in

improving the validity and generalizability of their study, as well as adopting a human-centered design and development approach.

Acknowledgments

We would like to thank Dr. Heath Yates for his help during the preparation of this book.

Leimin Tian

March 2021

Introduction to Applied Affective Computing

Affective computing is a nascent field that sits at the cross-section between artificial intelligence and social and behavioral science. It studies how human emotions are perceived and expressed, which informs the design of intelligent agents and systems that can either mimic this to improve their intelligence or incorporate such knowledge to effectively understand and communicate with their human collaborators. Affective computing research has enjoyed significant advancements recently and is facing a critical transformation from exploratory studies to real-world applications. This leads to an emerging research topic, which we name as *applied affective computing*. In this chapter, we will give our definition of applied affective computing and outline the unique challenges that differentiate applied affective computing from traditional affective computing research. We will also provide an overview of this book and its contributions.

1.1 Affective Computing

The concept of *affective computing* was first proposed more than two decades ago by Prof. Rosalind Picard in her paper [Picard 1995] and book [Picard 2000] published under the same name. She defined affective computing as “computing that relates to, arises from, or influences emotions.” Affective computing research can have a theory-focused perspective, aiming at understanding how people perceive and express emotions. For example, Marsella et al. [2010] reviewed computational models of emotions and how such models can be used to test a cognitive science theory, to implement artificial emotions in intelligent agents, or to establish a mental model of people interacting with an intelligent agent. Affective computing research can also have an application-focused perspective, aiming at designing intelligent systems that are capable of understanding and influencing people’s behaviors and/or mental states—for example, machine learning approaches to automatic emotion recognition [Swain et al. 2018, Sharma and

Dhall 2021], engineering approaches to synthesizing expressive emotional displays [Rosenthal-von der Pütten et al. 2018, Aneja et al. 2019], and human-computer/robot interaction studies involving social communications [Burgoon et al. 2017, Belpaeme et al. 2018].

Compared to when the concept of affective computing was first proposed in 1995, significant breakthroughs have been achieved in the field [Schuller and Schuller 2018]. For example, it is becoming increasingly common for data-driven automatic emotion recognition models to achieve human-level performance, especially for facial expression-based analysis [Sharma and Dhall 2021]. Beard et al. [2018] compared human and machine performance when classifying eight emotion categories, namely, anger, calm, disgust, fear, happiness, neutral, sadness, and surprise. They conducted experiments on three datasets, the CREMA-D dataset containing clips of actors portraying each emotion category as a short sentence with accompanying facial and vocal expressions [Cao et al. 2014], the RAVDESS dataset of actors portraying each emotion category as speech or song [Livingstone and Russo 2018], and the CMU-MOSEI dataset of YouTube opinion videos annotated with emotion labels [Zadeh et al. 2018]. Using a temporal dependent deep learning model and low-level features extracted from the raw signals on the CREMA-D dataset, the automatic emotion recognizer achieved higher accuracy than humans in audio-based emotion classification as well as in audio-visual emotion classification, as shown in Table 1.1.

In addition to improved performance in affective computing technologies, for instance, automatic emotion recognition, functional components based on affective computing studies have been continuously integrated into designs of interactive systems, for instance, in conversational agents [Poria et al. 2019] or social robots [Tsiourti et al. 2019]. In these affective computing systems, the intelligent agent interprets a user's behaviors and other information to inform its decisions in order to perform actions that achieve certain outcomes, as shown in Figure 1.1. For example, robots with the capability of interpreting patients' emotions and

Table 1.1 Human versus machine's accuracy in the emotion classification experiments of Beard et al. [2018]

Modality	Human	Machine
Audio	40.9%	41.5%
Visual	58.2%	52.5%
Audio-Visual	63.6%	65.0%

This table shows that automatic emotion recognition achieved higher accuracy than human performance for audio-based and audio-visual emotion recognition, but not for visual-based emotion recognition. Based on Beard et al. [2018].

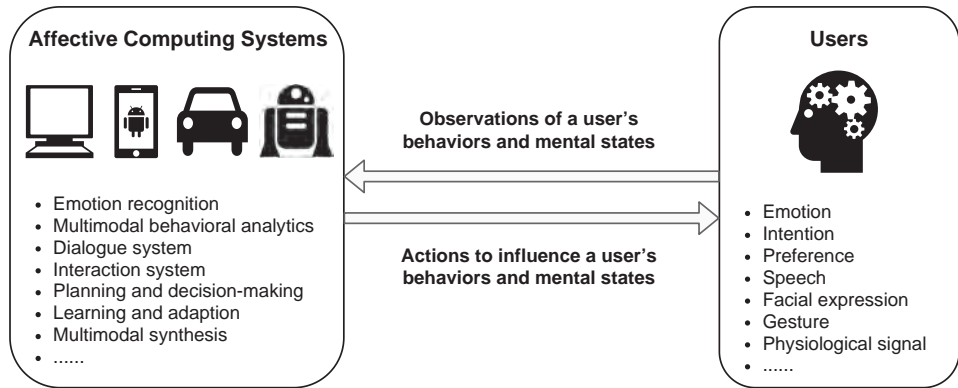


Figure 1.1 An affective computing system.

engaging in social interactions with patients have been used in psychosocial health interventions [Robinson et al. 2019] and have shown promising health benefits, especially in autism treatment [Ismail et al. 2019] and elderly care [Abdi et al. 2018].

Today, affective computing is rapidly growing as a research field and attracts both research and application interests from various disciplines, as shown in Figure 1.2. The Association for the Advancement of Affective Computing (AAAC, previously known as HUMAINE) was established for affective computing researchers worldwide, with its flagship biannual International Conference on Affective Computing and Intelligent Interaction (ACII) and the high impact *IEEE Transactions on Affective Computing* for researchers to share their work. Because of her invaluable contribution to the affective computing research community, Prof. Picard was awarded as one of the first AAAC Fellows at ACII 2019. With technical development rapidly maturing, we begin to see real-life applications of affective computing. For example, *Affectiva* [2020], a company founded in 2009, provides software that enables emotion analysis for clients in various domains, including the gaming industry, education, media, and healthcare. Thus, it is now the crucial moment to translate affective computing research advancement to intelligent systems that are capable of understanding and incorporating emotions in order to benefit individual users and society as a whole.

We capture this process of affective computing transformation from a controlled lab environment to in-the-wild application scenarios by proposing a novel research topic we call *applied affective computing*. We bring attention to the affective computing community to the pressuring challenges of applying research findings in real-life scenarios. In this book, we conduct extensive reviews on existing approaches for social signal processing, including *machine learning*, data collection, and signal processing. This provides references for researchers working on

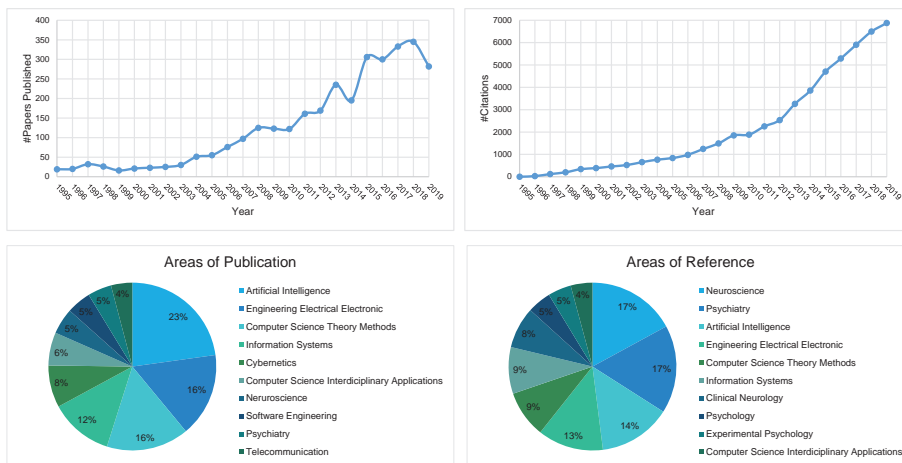


Figure 1.2 Publication data on affective computing according to Web of Science.

relevant topics on the state of the art and major challenges in the field. Moreover, we cover in-depth case studies in various application fields of affective computing, such as social robotics, healthcare, and education. We also summarize a set of guidelines for developing applied affective computing studies and systems for these different domains with a human-centered perspective. For affective computing researchers or researchers and practitioners who would like to apply affective computing to their fields, our book provides guidance for developing new generation human-centered artificial intelligence systems.

1.2 Applied Affective Computing

We define applied affective computing as research and application efforts that focus on understanding the effectiveness, outcomes, and personal and societal impact of affective computing technologies in real-life usage. Compared to affective computing, applied affective computing addresses the challenge of a holistic evaluation of computational models of emotions situated in context. In addition, applied affective computing maintains a human-centered perspective rather than a technology-centered perspective, and focuses on how affective computing systems influence their users' mental and physical states. The goal of applied affective computing research is to develop affective computing systems that are ethical and capable of achieving their intended outcomes in real-world applications emplaced in natural, unstructured, and complex environments.

Here we provide an example that illustrates our definition of applied affective computing: Automatic emotion recognition from speech has been a widely studied topic in affective computing. There has been significant improvement in

the performance of speech-based emotion recognition as evaluated on various datasets of emotional speech [Akçay and Oğuz 2020]. For applied affective computing research on speech emotion recognition, we shift the focus from identifying effective emotion recognition approaches benchmarked on a dataset to understanding how a model's performance may be influenced by contextual factors, for instance, user traits (e.g., age or personality), or type of speech (e.g., task-oriented dialogue or chit-chat). The main contextual factors are identified based on the analysis of the application goal of the speech emotion recognition model. In an example of a system aimed at identifying frustrated customers at a call center, the emotion recognition model would need to accommodate noises over the line among other factors. Moreover, beyond evaluating the accuracy of the speech emotion recognition model at identifying a frustrated customer from their calls, this model needs to be incorporated with other components of the system, for instance, a dialogue manager that selects the appropriate mitigation strategies or a knowledge base of available services specific to the call center. By assessing this integrated system on metrics, for example, the satisfaction of callers, the speech emotion recognition system's outcome is measured in a situated manner.

1.2.1 Challenges of Applied Affective Computing

Applied affective computing research is in its infancy. In order to develop an effective and ethical affective computing system, researchers face challenges that we will discuss below:

- Having a good understanding of the users and their expectations.
- Having functions that maintain their performance in varied environments.
- Having the capability to learn from accumulated observations.
- Having an integrated system that combines affect with other communicative and cognitive processes and with insights from related domains.
- Having a good understanding of the societal impacts of affective computing technologies to develop ethical research protocols and regulated applications.

1.2.1.1 Developing a Human-centered Affective Computing System

Affective computing systems address human emotion, cognition, and communication. Applied affective computing systems are often designed to monitor or influence the behaviors and/or mental states of their users. Therefore, they require a human-centered perspective during the design and evaluation processes

to address the needs and variances in the users and to adjust to the specific application scenarios [Norman and Draper 1986]. The new research area of behavioral analytics [Oviatt et al. 2018] has shed light on the recent development of human-centered multimodal interfaces and their applications in various domains, such as healthcare or education. To develop a human-centered affective computing system, the users should be involved in the design and evaluation processes so as to incorporate their expectations and needs. The system should address individual variances and preferences and be inclusive of different user traits. For example, emotional expressions generated by a system designed for dementia care purposes should utilize more pronounced, multimodal expressions to address the possible decline in emotional perception of dementia patients [Smiljanic and Gilbert 2017, Yatsuda et al. 2018]. To address individual variances and preferences, there have been increasing efforts in realizing personalization in interactive systems [Augstein et al. 2019, Irfan et al. 2019]. In Chapters 8 and 9, we provide case studies of applied affective computing for human-robot interaction and well-being to illustrate the design and evaluation of human-centered affective computing systems.

1.2.1.2 Developing a Reliable Affective Computing System

For an affective computing system to leave controlled lab environments and achieve its designed outcomes when applied in real-life settings, it is critical that its functionalities continue to deliver reliable performance. The topic of in-the-wild emotion recognition has received increasing interest from the affective computing community, which addresses one aspect of developing reliable affective computing systems [Picard 2011]. There have been benchmarking efforts, for example, the Emotion recognition in the wild challenges (EmotiW) [Dhall et al. 2020], in the research community to develop robust and accurate emotion recognition models that are resistant to noise and variance. To ensure the reliability of emotion recognition approaches, more rigorous evaluation methodologies are adopted in recent studies, for example, cross-corpora analysis [Eyben et al. 2015]. In addition, researchers have benefited from the advancement in crowd-sourcing platforms and vast online contents to generate high-fidelity datasets of in-the-wild affective interaction and experiences, for example, the CMU-MOSEI dataset [Zadeh et al. 2018]. To ensure the quality of data annotation, various measures are in place to account for subjectivity and uncertainty during emotion annotation and data processing and modeling, for instance, retaining distribution of labels for a data instance rather than one aggregated label per instance and modeling such distribution [Rizos and Schuller 2020]. In Chapters 3 and 4, we will provide a review on current machine learning and data processing approaches that benefit the

advancement of applied affective computing. Moreover, we offer case studies on in-the-wild emotion recognition in Chapter 6 to further discuss the reliability and robustness of affective computing systems.

1.2.1.3 Developing an Adaptive Affective Computing System

Many applied affective computing systems are designed to establish or maintain a social relationship with the users to achieve their intended goals. For example, intelligent tutoring systems able to build a good rapport with the tutees can lead to increased trust, learning experience, and learning outcomes [Zhao et al. 2016, Pan and Hamilton 2018]. To develop a social relationship with users, longitudinal interactions are required. Long-term interactions are also common in application scenarios, for instance, mental and physical healthcare, where a treatment program can contain multiple sessions over the span of months or even years. In such longitudinal interactions, it is necessary for an applied affective computing system to learn from and adapt to the shifting interaction dynamics and contexts. In addition, the capability to adapt and provide longitudinal services will help address attrition in use, which has been a major issue in commercialized affective computing systems, for example, social robots [Hoffman 2019]. More discussion and case studies on how reinforcement learning and other adaptation methods are adopted in applied affective computing can be found in Chapter 6.

1.2.1.4 Developing an Integrated Affective Computing System

In applied affective computing, emotion is not modeled in isolation. Instead, it is embedded in an integrated framework that includes other components that monitor or simulate mental states, cognition, and communicative functions. For instance, a social conversational agent requires dialogue capabilities, a knowledge base of potential topics, and the ability to address context, for instance, users' intentions, in addition to affective capabilities and social skills [Cohen 2018]. Similarly, the design of emotional expressions of a social robot needs to address the physical appearance and hardware capabilities of the robot [Song and Yamada 2017]. Applied affective computing is also an interdisciplinary research topic where collaborative efforts from domains including, but not limited to, psychology, AI, human-computer interaction (HCI), robotics, engineering, social science, and medical science are critical to developing meaningful hypotheses and solid methodologies. For example, therapists' insights are critical during the iterative design and improvement of mental well-being applications [Fitzgerald and McClelland 2017]. In Chapters 8, 9, and 10 we offer case studies of applied affective computing for social robots, for healthcare, and for built environments to illustrate the design and evaluation of integrated and interdisciplinary affective computing systems.

1.2.1.5 Developing an Ethical Affective Computing System

With the increasing application of affective computing technologies, there have been growing concerns on potential misuse of the technology and biases in the system [Zuboff 2019], for example, the ban on facial recognition and analysis technologies [Conger et al. 2019]. Thus, it is critical to understand the societal impact of applied affective computing and establish appropriate measures that ensure researchers are guided toward responsible research and practitioners are regulated toward ethical applications with societal benefits. As an emerging field with a short history, affective computing can learn from closely related fields that are more established, for example, HCI. The HCI field faces similar challenges in ensuring an ethical future of intelligent interactive systems, and have identified that collaboration between stakeholders, such as researchers, designers, users, and policymakers, is key to addressing the increasing concerns around fairness, accountability, and transparency of intelligent systems [Abdul et al. 2018]. In Chapter 11 we will further discuss ethical considerations and guidelines to responsible applied affective computing.

1.2.2 A Living Ontology of Affective Computing

Addressing the above challenges and facilitating the transformation from affective computing to applied affective computing require collaborative efforts. Therefore, during the 29th International Joint Conference on Artificial Intelligence and the 17th Pacific Rim International Conference on Artificial Intelligence (IJCAI-PRICAI 2020), as part of the 4th IJCAI workshop on Artificial Intelligence in Affective Computing (AffComp 2020) [Healey et al. 2020], we organized an open discussion on the scope and directions of affective computing. This hour-long discussion by workshop participants from various domains and both academia and industry backgrounds resulted in a mindmap of affective computing, as shown in Figure 1.3. To continuously develop a shared vocabulary of affective computing, we welcome contributions to a living concept board of this mindmap.¹

In this book, we will expand on the multifaceted concept of applied affective computing through a review of state-of-the-art approaches, as well as an in-depth discussion of case studies from different domains. To support a better understanding of affective computing that can be adopted in future studies or applications, as readers follow the book we invite them to create their version of this mindmap

1. Concept board of applied affective computing: <https://usulaep.conceptboard.com/board/fyic-x76t-ha3m-tt2n-442u>. History replay of how this ontology was developed during the workshop can be reviewed here: <https://usulaep.conceptboard.com/history/c03473b9-c18c-45c4-9147-25258fcb1d9f>.

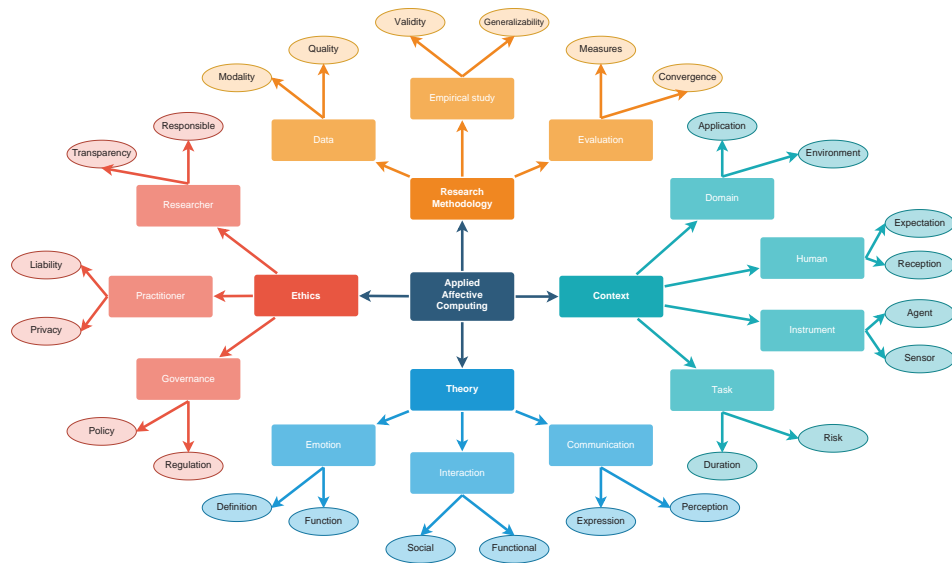


Figure 1.3 A mindmap of applied affective computing.

by growing the nodes with guidelines and challenges specific to their research or application context.

1.3 Book Overview and Highlights

In Chapters 2, 3, and 4 of this book we will provide reviews on state-of-the-art approaches to affective computing. After these comprehensive reviews, we will provide case studies in Chapters 5, 6, 7, 8, 9, and 10 to allow the readers to have a better understanding of the various aspects and challenges of applied affective computing. In Chapter 11, we further discuss the general issues of ethical affective computing. Finally, we summarize general guidelines and open questions of applied affective computing in Chapter 12. In this section, we offer an overview of the highlights in each chapter.

Chapter 2 reviews the theoretical background of emotion research. As affective computing researchers are often consumers of social science and behavioral science theories, it is important to base the design of an affective computing study on verified findings of human emotion and communication. What emotions are and how they function concerning other cognitive functions are actively researched in current psychology and cognitive science studies. For example, the evolutionary approach to emotion studies how it enables humans and other animals to generate adaptive behaviors in response to stimuli. Such theories can then be adopted to

design an adaptive intelligent agent capable of achieving its goals in a shifting environment. In addition, studies of emotion as a social-communicative function shed light on the design of interactive systems aimed at understanding and/or influencing human behaviors and mental states. Because of the context dependency in emotion perception and expression, affective computing researchers are also encouraged to tailor the emotion theory adopted to their specific study contexts, such as target user groups or task locations.

Chapter 3 gives an overview of *machine learning* approaches used in state-of-the-art affective computing studies with a focus on building an emotion-aware system. For supervised learning approaches, multimodal data of emotions are collected and labeled to allow machine learning models to infer the relationship between emotional labels and characteristics of the data. Among the various modalities that affect can be perceived and expressed: (1) studies in the visual modality often focus on facial expressions and body gestures, (2) studies in the audio modality often focus on speech intonation, (3) studies in the lexical modality often focus on speech transcription and written text, and (4) studies in the physiological modality often adopt data collected through wearable sensors that measure brain activities, heart activities, muscle activities, skin electrical conductivity, and gaze patterns. While identifying affective features in each of these modalities is crucial for developing accurate emotion recognition models, it is equally important to develop model structures that fuse these modalities effectively. Due to the limited availability of labeled emotional data, transfer learning approaches are commonly adopted to improve the performance of emotion recognition using knowledge from related domains, for example, sentiment analysis of lexical information. Current breakthroughs in deep neural networks also benefit affective computing by presenting end-to-end approaches to automatic emotion recognition and synthesis. In current machine learning studies of affective computing, beyond developing effective features and models for improving the accuracy of emotion recognition models, researchers are investing increasing efforts in incorporating contextual information in these models to improve their robustness.

Chapter 4 reviews the data collection and processing technologies used in existing emotion databases. As Chapter 3 illustrates, the quality of a computational model of emotions is directly dependent on the quality of the emotional data this model has been trained on. Therefore, affective computing researchers have focused on collecting high-quality datasets that are multimodal and naturalist in current studies. In these datasets, the audio and visual modalities are the most commonly recorded. In addition, physiological data have been a recent addition with wearable sensors becoming increasingly available. Compared to portrayed emotions often collected in earlier datasets, recent datasets have favored

naturalist data that are either collected by using stimuli verified by psychological studies or by gathering in-the-wild contents on social media platforms. When collecting data in naturalist settings, noises are inevitable, and data and feature engineering techniques, for example, feature normalization, can benefit the affective computing models to distill relevant information and to become more robust.

Chapter 5 provides a case study of emotion perception in the wild, which is the foundation of many affective computing systems. We investigated existing approaches to how emotions are represented in the wild, how data-driven deep learning models can be applied to learn the emotion representations, and how these models are evaluated. In particular, we highlighted the importance of contextual information for emotion recognition in the wild. For example, for emotion recognition models aimed at monitoring a person's mental health states to provide timely support, they need to take into account the person's mood and life events. Similarly, while many emotion recognition models are trained on data with a single participant exposed to emotional stimuli, in an application they may need to adapt to group interaction scenarios, for example, a cocktail party. Therefore, while it is essential to continue advancing computational approaches to improve the accuracy and robustness of emotion recognition, researchers are encouraged to explore learning-based approaches that allow a trained emotion recognition model to adapt to the complexity in real-world situations and to maintain its effectiveness in long-term deployment.

Chapter 6 extends on the learning-based approaches identified in Chapter 5 and focuses on case studies of reinforcement learning for affective computing. This includes both how artificial emotions can benefit an intelligent agent in adapting to the environment and how an interactive system can adopt user emotions as a reward signal to realize personalization through accumulated experiences. The key challenges for intelligent agents with artificial emotions are how they can improve sample efficiency, that is, being able to distill knowledge from fewer trials, and how they can transfer the knowledge acquired to similar tasks or environments.

Chapter 7 investigates the expression aspect of emotions and how an affective computing system can synthesize natural and believable emotional displays to convey its intentions in a human-understandable manner. As human communications are multimodal, it is also beneficial for an intelligent agent to express emotions through modalities, such as facial expressions, gaze patterns, and expressive speech synthesis. Because of the differences in the appearances of artificial agents, such as virtual characters or robots, simply mimicking human expressions in the design of artificial emotional expressions may either be less convincing or may fall into the “uncanny valley.” Thus, researchers may learn from principles

in performance art and animation design when generating synthesized emotional expressions and creating a sense of agency. Similar to emotion recognition, emotion expression is also context dependent, and expressive nuances, for instance, cultural differences, should be considered.

Chapter 8 provides a case study of affective computing in human–robot interaction and social robotics. In particular, it discusses challenges in incorporating social intelligence in a robot, and how errors observed in human–robot interaction can become opportunities for a robot to improve its social intelligence. An example study of human–robot collaboration demonstrates the benefits of artificial emotions and emotional expressions in allowing a robot to convey its intentions clearly and to motivate actions in human observers in alignment with the robot’s goals. Another example study exploring how the general public perceives social robots in public spaces highlights the importance of context-awareness in affective computing systems, moreover, the benefits of multidisciplinary collaborations in developing effective methodologies for understanding the societal impacts of such systems.

Chapter 9 focuses on affective computing applications for enhancing people’s physical and mental well-being. Affective computing research on well-being is closely related to the concept of positive computing and is centered around eight pro-social positive emotions, namely love, gratitude, hope, joy, forgiveness, compassion, faith, and awe. Because well-being is a long-term goal, affective computing systems designed to improve users’ well-being require the ability of longitudinal interaction and adaption enabled by data collected in a ubiquitous, non-invasive manner, as well as the incorporation of expert knowledge from healthcare professionals. Several example systems were reviewed, for example, mobile phone applications that help young adults regulate their emotional states in daily life. Beyond active sensing, passive sensing in a home environment and in vehicles provides a contextualized perception of users’ physical and mental states. Beyond sensing, an affective computing system needs to choose appropriate actions and interventions for realizing well-being-related goals, such as through conversations, recommendations, and behavioral modification.

Chapter 10 presents a case study of affective computing in built environments, which investigates people’s emotional perception and experience of buildings and landscapes. In an example study, how urban infrastructure, such as greenery area or street lighting, influences the physical and mental health of people was investigated. By leveraging machine learning approaches and cross-platform sensing capabilities, multimodal data on participants’ activities and affective responses were collected during their walk through the areas of interest. Compared to traditional analytic approaches, affective computing provides additional insights into

people's experience of a space, which can inform the planning of urban spaces that are both functional and enjoyable to the occupants.

Chapter 11 discusses ethical considerations of affective computing research and applications, including biases in data and models of affective computing, privacy concerns of invasive surveillance involved in some affective computing systems, consequences of a system capable of influencing a user's behaviors and mental states, and societal impacts of affective computing technologies. We also discuss approaches to fair, transparent, and responsible affective computing application.

Chapter 12 summarizes a set of guidelines for best practices in affective computing research and application. To illustrate how the guidelines and topics covered in this book can yield benefits, we provide a scenario-based discussion of how affective computing researchers and practitioners may acquire an overview of the field, refine their study design, and iterate their applications. In addition, we discuss the open challenge of improving the validity and generalizability of affective computing research methodologies. We also identify future directions of the field.

1.4 Contributions

Our book is not the first one on affective computing, and certainly not the last. The most closely related titles to this book are *The Oxford Handbook of Affective Computing* [Calvo et al. 2015], *Emotion-Oriented Systems: the HUMAINE Handbook* [Petta et al. 2011], *A Blueprint for Affective Computing: A Sourcebook and Manual* [Scherer et al. 2010], and *Handbook of Research on Synthetic Emotions and Sociable Robotics: New Applications in Affective Computing and Artificial Intelligence* [Vallverdú 2009].

Calvo et al. [2015] presents a thorough overview of affective computing, covering its theories, methodologies, and applications. It was designed to be a comprehensive textbook to introduce students and researchers to affective computing. It contains a collection of chapters by top researchers in the field on a wide scope of topics, and individual chapters can be read independently of other chapters to deepen a reader's understanding of that specific topic of interest. Similar to Calvo et al. [2015], Petta et al. [2011] is a collection of affective computing studies with a focus on presenting a broad overview of the field. It was an initiative by the HUMAINE community (currently known as AAAC) to bring together distinguished researchers from affective computing and other related fields to offer newcomers to the area a systematic introduction to affective computing. Compared to Calvo et al. [2015], Petta et al. [2011] focuses on emotions in interaction and communication. We recommend readers interested in obtaining a systematic understanding of the affective computing research area to read Calvo et al. [2015] and Petta et al. [2011]. Compared to Calvo et al. [2015] and Petta et al. [2011], our book focuses on

the challenges of incorporating affective computing research findings in real-world applications and complex interaction systems or intelligent agents. Our review of affective computing studies focuses on recent efforts, which provides additions to the studies presented in [Calvo et al. \[2015\]](#) and [Petta et al. \[2011\]](#).

[Scherer et al. \[2010\]](#) explores the multidisciplinary root of affective computing and grounds affective computing with psychology, affective neuroscience, and philosophy research. It also highlights the benefits of interdisciplinary collaboration between these fields and the computer science and engineering fields in advancing affective computing. [Scherer et al. \[2010\]](#) provides a rich source for affective computing researchers to improve the theoretical soundness of their study design and system design. Thus, it is a great extension to Chapter 2 of our book. Compared to [Scherer et al. \[2010\]](#), our book has a more computer science-focused perspective and addresses technical approaches that researchers and practitioners may adopt to improve their affective computing systems.

[Vallverdú \[2009\]](#) investigates one aspect of affective computing, namely incorporating emotion perception and expression in social robots. For readers interested in this particular topic, [Vallverdú \[2009\]](#) provides in-depth discussion of theories and approaches that they can adopt. Thus, [Vallverdú \[2009\]](#) is a great extension to Chapter 8 of our book. Compared to [Vallverdú \[2009\]](#), our book investigates a wider scope of affective computing research and applications, for instance, affective computing for understanding people's experience of urban spaces or affective computing for adaptive in-vehicle environments.

In summary, compared to these related work, *Applied Affective Computing* focuses on the challenges in the real-life application of affective computing. It provides an up-to-date review of affective computing research. In addition to providing a general introduction to applied affective computing, in-depth case studies in our book deepen the discussion and summarize tangible guidelines for conducting applied affective computing research in different domains. Thus, it benefits both researchers and practitioners new to the field and those more experienced. We hope that affective computing researchers and practitioners can benefit from our book by directing their attention to address the various challenges in conducting contextualized, in-the-wild research and evaluations. For researchers and practitioners in related fields, such as artificial intelligence, human-computer/robot interaction, psychology, and social science, we hope that this book will provide a comprehensive overview of affective computing and encourage conversations and collaborations across the disciplines. Major contributions of this book include:

- We propose the concept *applied affective computing*, which addresses the challenges in applying current advancements in affective computing research to real-life scenarios.
- We provide a thorough and up-to-date review on various aspects of affective computing research and application, painting the landscape of the field that inspires innovations.
- We provide detailed case studies of various applied affective computing domains to illustrate practices and potentials of applied affective computing.
- We provide comprehensive guidelines on conducting applied affective computing research and designing applied affective computing systems.

What is applied affective computing

1. Definition of applied affective computing: research and application efforts that focus on understanding the effectiveness, outcomes, and personal and societal impact of affective computing technologies in real-life usage.
2. Applied affective computing versus affective computing: applied affective computing performs a holistic evaluation of computational models relating to emotions situated in context. It focuses on how affective computing systems influence their users and the interaction context. It aims to develop affective computing systems that are ethical and capable of achieving their intended outcomes in natural, unstructured, and complex environments.



Emotions as Studied in Psychology and Cognitive Science

2.1 Theories of Emotion

The definition of human emotions remains an open question [Adolphs and Anderson 2018]. A commonly used categorization classifies existing emotion theories as evolutionary approach, appraisal approach, and constructionism, which emphasizes the biological givens of emotions, the cause of emotions, and the components of emotions, respectively [Niedenthal and Ric 2017]. In particular, the evolutionary approach and the appraisal approach have been commonly used in current affective computing research [Poria et al. 2017].

2.1.1 The Evolutionary Approach

The evolutionary approach to emotion theories considers emotions as products of evolution—for example, emotion theories that model disgust as the adaptive function for avoiding toxins [Tybur et al. 2013] or fear as the adaptive function for evading predators [Mobbs et al. 2015]. Emotion theories following the evolutionary approach attempt to identify a set of basic, distinct, and universal emotion categories, that is, the Basic Emotion Theory (BET). One representative is the Big-6 emotion categorization [Ekman et al. 1972] based on studies of human facial expressions, which identifies anger, disgust, fear, happiness, sadness, and surprise as the six basic emotions. Beyond the Big-6 emotion categorization, other sets of basic emotion categories have also been proposed. For example, Cordaro et al. [2018] proposed a set of 22 emotions based on analyses of emotional expressions across five cultures. In addition, the development of neuroscience and neuroimaging has provided novel approaches to identifying and verifying basic emotions through evidence such as brain activities captured by fMRI [Panksepp and Watt 2011, Celeghein et al. 2017]. In Figure 2.1, we illustrate overlaps between three

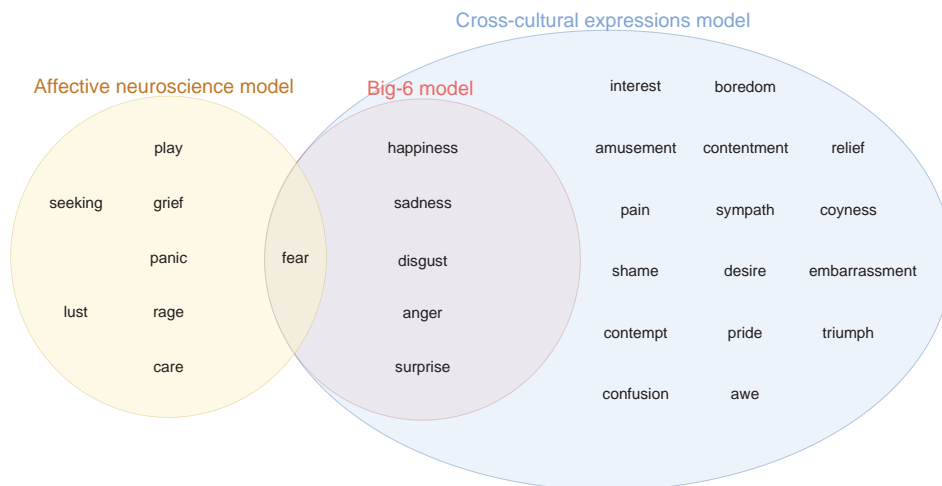


Figure 2.1 Overlaps between the Big-6 model [Based on Ekman et al. 1972], the cross-cultural expressions model [Based on Cordaro et al. 2018], and the affective neuroscience model [Based on Panksepp and Watt 2011].

common BET models, including the Big-6 model [Ekman et al. 1972], the cross-cultural expressions model [Cordaro et al. 2018], and the affective neuroscience model [Panksepp and Watt 2011].

In the recent progress of BET, there has been increasing attention on social functionalism approaches, which aim to explain the social functions of emotions and the fact that emotional experiences are personal and context-specific rather than universal, that is, “variation is the norm” [Barrett 2017]. In particular, the same event or stimulus can evoke a spectrum of emotional responses in different people and can be perceived differently, as demonstrated in existing studies of emotion induction and annotation tasks [Poria et al. 2017]. Behaviors and physiological responses across individuals also differ when they experience the same emotion. One example is fear, which was considered one of the basic emotions with an evolutionary root. A previous study identified that people can have varied actions in response to fear, such as freeze or flee, with different neural circuits involved [Gross and Canteras 2012]. Social functionalism BET argues that as humans are social animals, emotions evolved to address social activities and to serve as incentives for social behaviors, in addition to an individual’s need of surviving and adapting to the environment [Keltner et al. 2019]. The basic emotion categories are thus extended beyond the Big-6 to include interpersonal and social emotions such as sympathy or jealousy. In addition, emotional expressions developed as a result of coordination between individual and social behaviors, allowing people to provide

information to others regarding their internal states or to evoke specific responses in observers.

The subjective emotional experiences can be captured as a self-report of feelings. Existing studies following this approach have largely focused on verbal description-based approaches for measuring feelings,—for example, the Emotion Twenty Questions [Kazemzadeh et al. 2011] or the Positive and Negative Affect Schedule questionnaire [Watson and Clark 1999]. However, verbal-based measurement of feelings is less effective for non-verbal populations or across different languages and cultures. Measuring physiological or neural activation patterns provides an alternative to distinguishing different emotional feelings [Barrett 2017, Siegel et al. 2018]. As a limited understanding of the human brain may limit the emotion theories derived from neuroscience studies, affective neuroscience remains an active research area [Schore 2015, Watt 2018]. Furthermore, it is difficult to capture the functional aspects of emotions in relation to the environmental and temporal contexts with such physiological or neural approaches. Beyond individual variances, emotions are cultural-specific in their definition and expression. For instance, some languages do not make a clear distinction between anger and sadness [Russell 1991], while some emotional lexicons are “untranslatable” across languages [Ogarkova 2016]. In addition, a previous cross-cultural study found that norms in smiling, laughter, and expressions of positive emotions were related to the historical heterogeneity of the population in a country or a state [Niedenthal et al. 2018].

To balance between universality and variations, current BET research investigates complex emotion taxonomies placed in social and cultural contexts, and the boundaries and gradients of each emotion category [Cowen and Keltner 2017]. Although the Big-6 emotion categorization [Ekman et al. 1972] has been widely used in the current design of affective applications, researchers are encouraged to develop advanced emotion taxonomies that address the specific social and user context of the application to ensure its efficacy and to allow personalization and adaptation.

2.1.2 The Appraisal Approach

The appraisal approach to emotion theories focuses on defining and distinguishing emotions through cause–effect relationships. These emotion theories consider emotions as reactions to stimuli in the environment and argue that emotions are products of factual and evaluative cognitions. Emotion theories following the appraisal approach attempt to identify a set of dimensions from which span complex and compound emotions. With a focus on the causes of emotions, appraisal theories address the differentiated nature of emotional responses, such

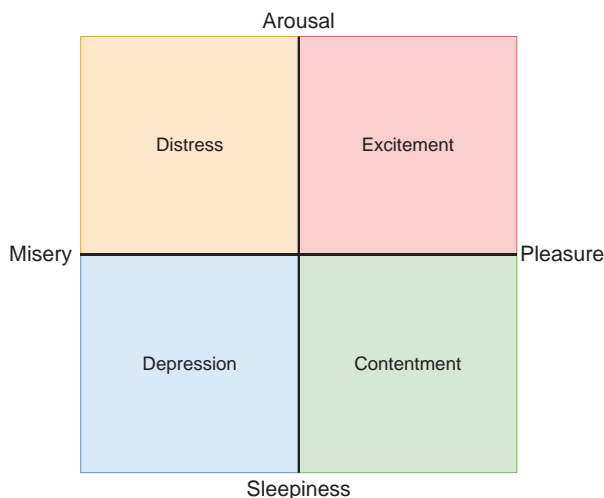


Figure 2.2 The circumplex model of emotion. Based on [Russell \[1980\]](#).

as individual, temporal, and cultural variations in emotional responses to the same stimuli, and how the same emotion can be evoked by a range of events and scenarios [[Scherer et al. 2001](#)].

One representative of appraisal-based emotion theory is the circumplex model of emotion [[Russell 1980](#)], which describes emotions with two dimensions, namely Arousal (level of excitement caused by a stimulus) and Valence (level of liking toward a stimulus) (Figure 2.2). Other than arousal and valence, the power or dominance dimension is often added to the emotional space to describe the level of control a person feels [[Marsella et al. 2010](#)]. These dimensions capture a one-to-one relationship between an appraisal combination (evaluation of events) that evokes the same emotion, invariant to specific events or situations [[Smith and Lazarus 1990](#)].

Beyond the circumplex model of emotion, other appraisal-based emotion theories with different sets of dimensions have been proposed. For example, [Adolphs and Andler \[2018\]](#) identified scalability (intensity, arousal), valence (pleasantness/unpleasantness, appetitive/aversive), persistence (temporal scale), generalization (generalized over stimuli and behavior), global coordination (engaging the whole organism), automaticity (having priority on behavioral control and requires volitional effort to regulate), and social communication (social-communicative signals) as fundamental characteristics of emotions. They argued that these attributes are capable of distinguishing emotion from other cognitive processes and can individualize different emotions.

A major constraint of appraisal emotion theories is that they assume emotions are temporally dependent and are elicited by current appraisals. Thus, appraisal

emotion theories often overlook long-term factors and their influence on short-term emotional changes such as moods or personality traits [Marsella and Gratch 2009]. The relation between emotions and stimuli has a complex “fan-in, fan-out” structure, that is, a stimulus can cause multiple emotions and the same emotion can be induced by various stimuli [Tappolet 2016]. Therefore, emotion theory alone is not sufficient to explain the causal relationship between stimuli and behaviors. Emotions interact with all other cognitive states. Thus, other cognitive processes, such as memory or attention, need to be included for a system-level view [Pessoa 2018].

The appraisal emotion theories are commonly used by affective computing researchers to design subtle or compound emotional expressions in response to specific events during interaction [Bartneck and Lyons 2009]. When defining emotions in the dimensional space, the emotion values can either be annotated as continuous real values, as discrete classes (e.g., high/low), or on Likert scales (e.g., 1 to 5 integer scores). In order to address individual variations in emotion perception, recent work has also employed ordinal rankings instead of absolute values to annotate emotions in the dimensional space [Yannakakis et al. 2017].

2.1.3 The Constructionism Approach

The constructionism approach to emotions assumes that emotions are created in the mind of a perceiver and are constructed of core affective changes situated with conceptual knowledge and exteroceptive sensation, as shown in Figure 2.3 [Lindquist 2013]. Such situated conceptualization allows researchers to take social and individual variances into account and analyze emotions from a social–psychological perspective.

Following the constructionism approach, researchers attempt to identify core affects and how the situated conceptualization process unfolds. In particular, core

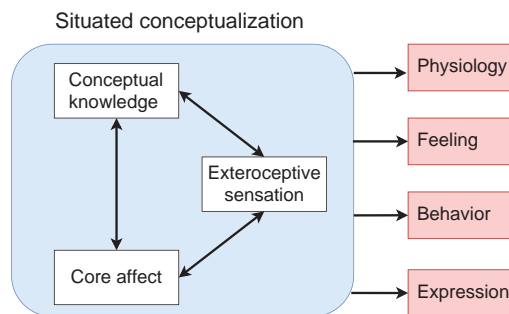


Figure 2.3 The constructionism approach to emotions. Based on Lindquist [2013].

affect can be described as a combination of interoceptive information, including homeostatic status, muscle activities, nervous system, neurochemical system, and hormonal system [Barrett and Bliss-Moreau 2009]. While interoceptive information represents the internal state of a person, exteroceptive sensation represents information from the external environment such as vision, audition, tactile, or olfactory sensations [Barrett 2009]. Specific emotional experiences or perceptions of oneself or other individuals can form conceptual knowledge, which shapes our understanding and predictions of new sensations and actions [Vigliocco et al. 2009].

To study the social and cultural perspective of emotions, researchers following the constructionism approach are especially interested in understanding how language interplays with emotional perception and experiences [Lindquist et al. 2015]. Studies in developmental psychology have shown that language allows the acquisition of emotional concepts in infants and children [Widen and Russell 2008, Halberstadt and Lozada 2011]. People continue to update and acquire conceptual knowledge throughout their lives, and language was also shown to help adults learn novel emotion concepts [Fugate et al. 2010]. Beyond knowledge acquisition, language can influence the situated conceptualization of emotions. For example, multilingual speakers are reported to have different emotional reactivity toward affective words in their dominant and non-dominant languages, or when recalling emotional events in different languages [Pavlenko 2014].

In affective computing research, the constructionism approach to emotion enables computational modeling of context, which can be combined with appraisal-based emotion models to enable context-aware inference [Ortony and Clore 2015]. Similar to the evolutionary and appraisal approaches, the constructionism approach to emotion is subject to its own limitations. In particular, it remains an open question whether the emotional context should be drawn at an individual's level or a collective level [Lyon 1995, Beatty 2010].

2.2 Emotion as an Adaptive Function

Psychology and neuroscience studies identified three main intrapersonal functions of emotions:

1. the attention-directing function, that is, emotions allow better allocation of limited cognitive resources [Simon 1967];
2. the informational function, that is, emotions provide high-level evaluation of events based on their relationship with the internal beliefs and desires [Slovic et al. 2005];

3. the motivational function, that is, emotions serve as a reward function that motivates us to strategize our behaviors based on emotional feedback [Oatley and Johnson-Laird 1987].

Following the basic emotion theories discussed in Section 2.1.1, emotions emerged to allow the most efficacious behavior in a given biological or social environment and enables flexible responses to complex or subtle living conditions [Lazarus and Lazarus 1991]. Following the appraisal emotion theories discussed in Section 2.1.2, people perceive the environment in relation to their goals as motivational relevance and congruence [Smith and Ellsworth 1985], and emotions reveal the appropriate actions to take given such evaluations. Following the social constructionism emotion theories discussed in Section 2.1.3, emotions lead to learning perceptions and actions that are considered normative in a given sociocultural context [Sarbin 2014].

A concept often discussed together with the adaptive function of emotion is coping, which refers to efforts to manage a threatening or harmful situation to reduce its adverse impact on the person [Smith and Kirby 2011]. In particular, coping leads to behaviors that either change the situation to bring it closer to one's goals or modifies one's goals to better align with the situation. When a situation is beyond one's ability to cope with it, psychological stress may be induced [Folkman and Lazarus 1984]. Thus, beyond strategizing behaviors, the adaptive functions of emotion are also critical to a person's well-being. The relationship between emotions and physical and mental well-being will be further discussed in Chapter 9.

The intrapersonal functions of emotion motivated affective computing researchers to implement artificial emotions in a robot or a virtual agent to improve its ability to generate adaptive behaviors. For example, the belief-desire-intention agent framework often integrates emotions to increase the believability of simulations or the efficiency of an agent [Adam and Gaudou 2016]. Furthermore, a computational model of emotion can benefit reinforcement learning agents in their decision-making [Moerland et al. 2018]. The use of emotions in reinforcement learning will be further discussed in Chapter 6.

2.3 Emotion as a Social Communicative Function

During interpersonal communications, emotions have been identified as cues that allow people to observe and perceive the latent cognitive states of their interaction partners and to express their own internal states [Vinciarelli et al. 2012]. Emotional expressions and reactions to other's emotions were found to be essential for establishing social relationships [Lewis 2008]. Emotional experiences navigate us through the complexities of the social world [FeldmanHall and Chang 2018].

In particular, while monitoring our progress toward achieving a specific social goal (e.g., the desire to affiliate with someone), negative emotions can function as an error signal that motivates us to adjust our decision policies and action strategies during the social interaction, while positive emotions can function as a reward signal and social incentive. In this section, we review emotions in the social-communicative context. In particular, we investigate the function of emotions in empathy, social intelligence, and social norms.

2.3.1 Empathy

One particular interpersonal emotional phenomenon that has attracted growing attention in recent affective computing studies is the computational analysis and simulation of empathy [Xiao et al. 2016]. Similar to emotions, various definitions of empathy have been adopted in current research on interpersonal communications [Batson 2009]. Cuff et al. [2016] reviewed 43 distinct definitions of empathy across various viewpoints and summarized empathy as an emotional response that “is similar to one’s perception and understanding of the stimulus emotion, with the recognition that the source of the emotion is not one’s own.” Decety and Meyer [2008] also found that the shared anchor of different definitions of empathy is the matching relation between the emotions of two people.

Empathy allows the development of shared representations, engagement, and simulation of the subjectivity of others. It is also important for the formation and maintenance of social relationships [Decety and Jackson 2004]. It often leads to behavioral and relational outcomes. For example, in psychotherapy sessions, the empathy of a therapist is a key behavioral marker for the quality and constructive outcomes of the therapy [Burke et al. 2004]. On the contrary, a lack of empathy is a major trait for socially aversive personality disorders [Paulhus and Williams 2002, Furnham et al. 2013] such as psychopathy [Hare 2003]. Therefore, an applied affective computing system with empathic functions leads to socially acceptable interaction designs and is required for establishing social relationships between the system and its users. An example is the development of affective grounding [Jung 2017] in current research, which refers to the coordination of affects in interaction with the purpose of building shared understanding on how behaviors should be interpreted emotionally and what counts as “appropriate.”

From the cognitive science perspective, empathy can be decoupled into two processes: an automatic process that captures the feelings of others (i.e., emotional contagion) and a controlled process that interprets those feelings (i.e., empathic understanding) [Heyes 2018]. The automatic process develops early in humans and has been found in a wide range of other animals [de Waal and Preston 2017], while the controlled process develops later and may be unique to humans. Previous work suggested that empathy is an associative learning process that is implemented by

the mirror neurons and constructed through social interaction [Gallese et al. 2004, Cook et al. 2014]. These learned matchings can be enhanced or broken by new experiences [Heyes 2018]. From the social science perspective, empathy promotes life-enhancing personal relationships, prosocial behaviors, and social justice [Batson 2011, Zaki 2018]. However, empathy might lead to biased perceptions and favoritism at the same time [Bloom 2017].

Other than empathy, the matching of emotion between people has also been studied in emotional contagion. Emotional contagion is “the tendency to automatically mimic and synchronize expressions, vocalizations, postures, and movements with those of another person’s and, consequently, to converge emotionally” [Hatfield et al. 1993]. Compared to empathy, emotional contagion focuses on the convergence of verbal and non-verbal behaviors in interpersonal communications, or entrainment [Narayanan and Georgiou 2013]. Co-occurrence in behavioral patterns of emotional expressions between conversational dyads has been shown to be a predictor of rapport in collaborative problem solving [Madaio et al. 2017]. The degree of convergence in human dyadic interaction has also been found to be positively correlated with the overall success of collaboration and positive attitudes of the participants [Levitan et al. 2012]. In educational scenarios, entrainment has been identified as an indicator for better learning outcomes and mutual understanding [Oviatt 2013].

In affective systems, empathy and entrainment are both implemented as an interactive strategy as well as a behavioral indicator of the users. For example, lexical and gestural entrainment has been observed in various human–robot interaction studies and can serve as an assessment of a user’s empathy and trust toward a robot, or to evoke certain behaviors in the users [Breazeal 2002a, Iio et al. 2015, Beňuš et al. 2018].

2.3.2 Social Intelligence and Socio-affective Competence

The ability to successfully conduct social interactions, that is, socio-affective competence, depends on the awareness and identification of social-emotional cues, the ability to process such cues, and the ability to decide on and express a normative response to these cues [Eisenberg 2001, Halberstadt et al. 2001]. One type of socio-affective competence studied extensively in human social communication is the Theory of Mind (ToM), which includes Cognitive ToM, referring to the ability to make inferences about the beliefs and motivations of other people, and Affective ToM, referring to the ability to make inferences about the emotions and feelings of other people [Premack and Woodruff 1978, Kalbe et al. 2010, Baksh et al. 2018]. A person’s performance in ToM tasks is often measured to evaluate their capability in social interactions, especially in the mental health context.