

AAAI Spring Symposium (SS-18)

Palo Alto, California, USA
26-28 March 2018

ISBN: 978-1-7138-2742-9

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2018) by Association for the Advancement of Artificial Intelligence
All rights reserved.

Printed with permission by Curran Associates, Inc. (2021)

For permission requests, please contact Association for the Advancement of Artificial Intelligence
at the address below.

Association for the Advancement of Artificial Intelligence
2275 East Bayshore Road
Suite 160
Palo Alto, California 94303
USA

Phone: 1-650-328-3123
Fax: 1-650-321-4457

<https://aaai.org/Press/press.php>

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

TABLE OF CONTENTS

AI AND SOCIETY: ETHICS, SAFETY AND TRUSTWORTHINESS IN INTELLIGENT AGENTS

From GOODBOT to BESTBOT	2
<i>Oliver Bendel</i>	
The Uncanny Return of Physiognomy	10
<i>Oliver Bendel</i>	
Maintaining the Humanity of Our Models	18
<i>Umang Bhatt</i>	
The Heart of the Matter: Patient Autonomy as a Model for the Wellbeing of Technology Users	23
<i>Emanuelle Burton, Kristel Clayville, Judy Goldsmith, Nicholas Mattei</i>	
Trustworthiness and Safety for Intelligent Ethical Logical Agents via Interval Temporal Logic and Runtime Self-Checking	29
<i>Stefania Costantini, Giovanni De Gasperis, Abeer Dyoub, Valentina Pitoni</i>	
Ethical Considerations for AI Researchers.....	37
<i>Kyle Dent</i>	
Interactive Agent that Understands the User	41
<i>Piotr Gmytrasiewicz, George Moe, Adolfo Morena</i>	
Toward Beneficial Human-Level AI... and Beyond.....	48
<i>Philip C. Jackson Jr.</i>	
Preferences and Ethical Principles in Decision Making.....	54
<i>Andrea Loreggia, Nicholas Mattei, Francesca Rossi, K. Brent Venable</i>	
Importance of Contextual Knowledge in Artificial Moral Agents Development.....	61
<i>Rafal Rzepka, Kenji Araki</i>	
Towards Provably Moral AI Agents in Bottom-Up Learning Frameworks	69
<i>Nolan P. Shaw, Andreas Stöckel, Ryan W. Orr, Thomas F. Lidbetter, Robin Cohen</i>	
Ethics as Aesthetic for Artificial General Intelligence	76
<i>Dan Ventura</i>	
An Architecture for a Military AI System with Ethical Rules.....	81
<i>Yetian Wang, Daniel Friyia, Kanzhe Liu, Robin Cohen</i>	
Architecting a Human-Like Emotion-Driven Conscious Moral Mind for Value Alignment and AGI Safety.....	88
<i>Mark R. Waser, David J. Kelley</i>	
Trustworthy Automated Essay Scoring without Explicit Construct Validity.....	95
<i>Patti West-Smith, Stephanie Butler, Elijah Mayfield</i>	
The Potential Social Impact of the Artificial Intelligence Divide	103
<i>Andrew B. Williams</i>	

ARTIFICIAL INTELLIGENCE FOR THE INTERNET OF EVERYTHING

Compositional Models for the Internet of Everything	107
<i>Spencer Breiner, Ram D. Sriram, Eswaran Subrahmanian</i>	
Internet of Things: Securing the Identity by Analyzing Ecosystem Models of Devices and Organizations.....	111
<i>Kai Chih Chang, Razieh Nokhbeh Zaeem, K. Suzanne Barber</i>	
Meta-Agents: Managing Dynamism in the Internet of Things (IoT) with Multi-agent Networks	117
<i>Hesham Fouad, Ira S. Moskowitz</i>	
Message Validation Pipeline for Agents of the Internet of Everything	119
<i>Boris Galitsky</i>	
Policy Issues Regarding Implementations of Cyber Attack Resilience Solutions for Cyber Physical Systems.....	128
<i>Barry M. Horowitz</i>	
On Stream-Centric Learning for Internet of Battlefield Things	135
<i>Brian A. Jalaiian, Alec Koppel, Andre Harrison, James Michaelis, Stephen Russell</i>	
Challenges and Characteristics of Intelligent Autonomy for Internet of Battle Things in Highly Adversarial Environments	145
<i>Alexander Kott</i>	
Artificial Intelligence for the Internet of Everything.....	152
<i>W. F. Lawless, Ranjeev Mittu, Donald Sofge</i>	
Active Inference in Multi-Agent Systems: Context-Driven Collaboration and Decentralized Purpose-Driven Team Adaptation	157
<i>Georgiy Levchuk, Krishna Pattipati, Adam Fouse, Robert McCormack, Daniel Serfaty</i>	
Viewing Machines as Teammates: A Qualitative Study	166
<i>Joseph B. Lyons, Sean Mahoney, Kevin T. Wynne, Mark A. Roebke</i>	
Valuable Information and the Internet of Things	171
<i>Ira S. Moskowitz, Stephen Russell</i>	
AI Enabled Blockchain Smart Contracts: Cyber Resilient Energy Infrastructure and IoT.....	179
<i>Michael Mylrea</i>	
The Smart Data Layer	185
<i>Magnus Sahlgren, Erik Ylipää, Barry Brown, Karey Helms, Airi Lampinen, Donald McMillan, Jussi Karlgren</i>	
The Web of Smart Entities — Towards a Theory of the Next Generation of the Internet of Things	189
<i>Michael Wollowski, John McDonald, Vishal Kapashi, Benjamin Chodroff</i>	

BEYOND MACHINE INTELLIGENCE: UNDERSTANDING COGNITIVE BIAS AND HUMANITY FOR WELL-BEING AI

Early Dementia Detection through Conversations to Virtual Personal Assistant.....	198
<i>Saleh Ahmed, Mahboob Qaosar, Rizka Wakhidatus Sholikah, Yasuhiko Morimoto</i>	

Measuring Cognitive Bias in Spoken Interaction and Conversation:	204
<i>Christina Alexandris</i>	
A Study on the UI of Musical Performance System and Score Representation	207
<i>Sachiko Deguchi</i>	
A Dynamic Learning Model for a Better Personalized Healthcare Using Mobile Health Tools.....	212
<i>Amy Wenxuan Ding</i>	
Customers' Retention Requires an Explainability Feature in Machine Learning Systems They Use	214
<i>Boris Galitsky</i>	
Retrieval System for Data Utilization Knowledge Integrating Stakeholders' Interests	221
<i>Teruaki Hayashi, Yukio Ohsawa</i>	
Policy Decision Support System in Aging Society Based on Probabilistic Latent Spatial Semantic Structure Modeling.....	223
<i>Ayae Ide, Yoichi Motomura, Takao Terano</i>	
Texture Suggestion System Considering the Elderly's Preference on 3D Printing	230
<i>Tomohiko Inazumi, Jinhwan Kwon, Shinsaku Hiura, Maki Sakamoto</i>	
The Challenges for Understanding Cognitive Bias and Humanity for Well-Being AI — Beyond Machine Intelligence	237
<i>Takashi Kido, Keiki Takadama</i>	
Developing a Dataset for Personal Attacks and Other Indicators of Biases	239
<i>John Licato, Mark Boger, Zhitian Zhang</i>	
A Personalized Method for Calorie Consumption Assessment.....	247
<i>Yunshi Liu, Pujana Paliyawan, Takahiro Kusano, Tomohiro Harada, Ruck Thawonmas</i>	
IoT-based Emotion Recognition Robot to Enhance Sense of Community in Nursing Home	253
<i>Shintaro Nagama, Masayuki Numao</i>	
Active Online Learning Architecture for Multimodal Sensor-based ADL Recognition	259
<i>Nobuyuki Oishi, Masayuki Numao</i>	
Estimation of Personalized Value through the Analysis of Conversational Data Assisted by Coimagination Method.....	267
<i>Mihoko Otake, Masato S. Abe, Masahiro Nochi, Eij Shimizu</i>	
Hybrid Sensing and Wearable Smart Device for Health Monitoring and Medication: Opportunities and Challenges	269
<i>Mahboob Qaosar, Saleh Ahmed, Chen Li, Yasuhiko Morimoto</i>	
From Algorithms to Heuristics: Will Androids Ever Make Freudian Slips?.....	275
<i>Sadeq Rahimi</i>	
Sleep Stage Re-Estimation Method According To Sleep Cycle Change.....	280
<i>Yusuke Tajima, Akinori Murata, Tomohiro Harada, Keiki Takadama</i>	
Can Machine Learning Correct Commonly Accepted Knowledge and Provide Understandable Knowledge in Care Support Domain? Tackling Cognitive Bias and Humanity from Machine Learning Perspective	285
<i>Keiki Takadama</i>	

Study of Analytical Methods on the Relationship between Sleep Quality and Stress with a focus on Human Circadian Rhythm.....	291
<i>Ryo Takano, Satoshi Hasegawa, Yuta Umenai, Takato Tatsumi, Keiki Takadama, Toru Shimuta, Toru Yabe, Hideo Matsumoto</i>	
Improving Sleep Stage Estimation Accuracy by Circadian Rhythm Extracted from a Low Frequency Component of Heart Rate	297
<i>Akari Tobaru, Fumito Uwano, Takuya Iwase, Kazuma Matsumoto, Ryo Takano, Yusuke Tajima, Yuta Umenai, Keiki Takadama</i>	
Ensemble Heart Rate Extraction Method for Biological Data from Water Pressure Sensor	304
<i>Fumito Uwano, Keiki Takadama</i>	
Does Digital Dementia Exist?	310
<i>Hideya Yamamoto, Kaoru Ito, Chihiro Honda, Eiji Aramaki</i>	

DATA EFFICIENT REINFORCEMENT LEARNING

Efficient Exploration for Constrained MDPs	313
<i>Majid Alkaee Taleghan, Thomas G. Dietterich</i>	
Towards a Data Efficient Off-Policy Policy Gradient.....	320
<i>Josiah P. Hanna, Peter Stone</i>	
Bayesian Q-learning with Assumed Density Filtering	324
<i>Heejin Jeong, Daniel D. Lee</i>	
State Abstraction Synthesis for Discrete Models of Continuous Domains.....	331
<i>Jacob Menashe, Peter Stone</i>	
Run, Skeleton, Run: Skeletal Model in a Physics-Based Simulation	338
<i>Mikhail Pavlov, Sergey Kolesnikov, Sergey Plis</i>	
Inverse Reinforcement Learning via Nonparametric Subgoal Modeling	344
<i>Adrian Šošić, Abdelhak M. Zoubir, Heinz Koepl</i>	
Multiagent Soft Q-Learning	351
<i>Ermo Wei, Drew Wicke, David Freelan, Sean Luke</i>	
Hierarchical Approaches for Reinforcement Learning in Parameterized Action Space.....	358
<i>Ermo Wei, Drew Wicke, Sean Luke</i>	

THE DESIGN OF THE USER EXPERIENCE FOR ARTIFICIAL INTELLIGENCE (THE UX OF AI)

Usability Issues and Guidance for Flexible Execution of Procedural Work	366
<i>Dorrit Billman, Debra Schreckenghost</i>	
Building Bridges: A Case Study in Structuring Human-ML Training Interactions via UX.....	370
<i>Johanne Christensen, Benjamin Watson, A. J. Rindos, Stacy Joines</i>	
User Interfaces and Scheduling and Planning: Workshop Summary and Proposed Challenges	373
<i>Richard G. Freedman, Tathagata Chakraborti, Kartik Talamadupula, Daniele Magazzeni, Jeremy D. Frank</i>	

Designing for Trust with Machine Learning	378
<i>Fabien Girardin, Pablo Fleurquin</i>	
Revealing Actionable Simplicity in Data	382
<i>Nick Gisolfi, Artur Dubrawski</i>	
Artificial Digitality	386
<i>Kuldeep Gohel</i>	
Towards Natural Cognitive System Training Interactions: A Preliminary Framework	388
<i>Erik Harpstead, Christopher J. MacLellan, Robert P. Marinier, Kenneth R. Koedinger</i>	
Design Methods to Investigate User Experiences of Artificial Intelligence	394
<i>Karey Helms, Barry Brown, Magnus Sahlgren, Airi Lampinen</i>	
Designing Therapeutic Care Experiences with AI in Mind	399
<i>Aisling Kelliher, Barbara Barry</i>	
Insectile Indices	403
<i>Yeawon Kim</i>	
The Importance of UX for Machine Teaching	407
<i>Martin Lindvall, Jesper Molin, Jonas Löwgren</i>	
Trees of Knowledge: Designing with Artificial Intelligence in the Urban Landscape	411
<i>Xiaoxuan Liu, Godiva Veliganilao Reisenbichler</i>	
The UX of AI: Using Google Clips to Understand how a Human-Centered Design Process Elevates Artificial Intelligence	414
<i>Josh Lovejoy</i>	
FutureCrafting. A Speculative Method for an Imaginative AI	419
<i>Betti Marenko</i>	
A Panel on Cybernetics and the User Experience of AI Systems	423
<i>Nikolas Martelaro, Wendy Ju</i>	
Knowledge Design — Towards an Inclusive, AI Design Practice	425
<i>Christine Meinders, Selwa Sweidan</i>	
Challenges and Methods in Design of Domain-specific Voice Assistants	431
<i>Sarah Mennicken, Ruth Brillman, Jennifer Thom, Henriette Cramer</i>	
Intelligent Devices Retirement Preserve: (un) Natural Wonders	436
<i>Michael Milano</i>	
Talk to Me About Pong: On Using Conversational Interfaces for Mixed-Initiative Game Design	439
<i>Afshin Mobramaein, Jim Whitehead, Chandranil Chakrabortii</i>	
How Can I Cook with This: User Experience Challenges for AI in the Home Kitchen	443
<i>Johnathan Pagnutti</i>	
Procedure Automation: Sharing Work with Users	446
<i>Debra Schreckenghost, Scott Bell, David Kortenkamp, James Kramer</i>	
Assessing and Addressing Algorithmic Bias - But Before We Get There	450
<i>Aaron Springer, Jean Garcia-Gathright, Henriette Cramer</i>	

What Are You Hiding? Algorithmic Transparency and User Perceptions	455
<i>Aaron Springer, Steve Whittaker</i>	
Hey Scout: Designing a Browser-Based Voice Assistant.....	460
<i>Janice Y. Tsai, Jofish Kaye</i>	
Committee of Infrastructure: Civic Agency and Representation.....	463
<i>Jason Wong</i>	
Machine Learning as a UX Design Material: How Can We Imagine Beyond Automation, Recommenders, and Reminders?	467
<i>Qian Yang</i>	

INTEGRATING REPRESENTATION, REASONING, LEARNING, AND EXECUTION FOR GOAL DIRECTED AUTONOMY

Robot Behavioral Exploration and Multi-modal Perception using Dynamically Constructed Controllers.....	474
<i>Saeid Amiri, Suhua Wei, Shiqi Zhang, Jivko Sinapov, Jesse Thomason, Peter Stone</i>	
Validation of Hierarchical Plans via Parsing of Attribute Grammars	482
<i>Roman Bartak, Adrien Maillard, Rafael C. Cardoso</i>	
Situated Planning for Execution Under Temporal Constraints.....	490
<i>Michael Cashmore, Andrew Coles, Bence Cserna, Erez Karpas, Daniele Magazzeni, Wheeler Ruml</i>	
Creating and Using Tools in a Hybrid Cognitive Architecture	497
<i>Dongkyu Choi, Pat Langley, Son Thanh To</i>	
Flexible Goal-Directed Agents' Behavior via DALI MASs and ASP Modules	505
<i>Stefania Costantini, Giovanni De Gasperis</i>	
Perspectives on the Validation and Verification of Machine Learning Systems in the Context of Highly Automated Vehicles	512
<i>Werner Damm, Martin Fränzle, Sebastian Gerwinn, Paul Kröger</i>	
SiRoK: Situated Robot Knowledge - Understanding the Balance Between Situated Knowledge and Variability.....	516
<i>Angel Andres Daruna, Vivian Chu, Weiyu Liu, Meera Hahn, Priyanka Khante, Sonia Chernova, Andrea Thomaz</i>	
Teaching Virtual Agents to Perform Complex Spatial-Temporal Activities.....	525
<i>Tuan Do, Nikhil Krishnaswamy, James Pustejovsky</i>	
Planning Hierarchies and their Connections to Language.....	532
<i>Nakul Gopalan</i>	
Learning Generalized Reactive Policies using Deep Neural Networks.....	537
<i>Edward Groshev, Aviv Tamar, Maxwell Goldstein, Siddharth Srivastava, Pieter Abbeel</i>	
Constraint-Based Online Transformation of Abstract Plans into Executable Robot Actions.....	549
<i>Till Hofmann, Victor Mataré, Stefan Schiffer, Alexander Ferrein, Gerhard Lakemeyer</i>	
Learning to Act in Partially Structured Dynamic Environment	554
<i>Chen Huang, Lantao Liu, Gaurav Sukhatme</i>	

Representation, Use, and Acquisition of Affordances in Cognitive Systems.....	560
<i>Pat Langley, Mohan Sridharan, Ben Meadows</i>	
Learning Planning Operators from Episodic Traces.....	564
<i>David Menager, Dongkyu Choi, Mark Roberts, David W. Aha</i>	
Human-Agent Teaming as a Common Problem for Goal Reasoning.....	572
<i>Matthew Molineaux, Michael W. Floyd, Dustin Dannenhauer, David W. Aha</i>	
Interaction and Learning in a Humanoid Robot Magic Performance.....	578
<i>Kyle John Morris, John Anderson, Meng Cheng Lau, Jacky Baltes</i>	
Position Paper: Reasoning About Domains with PDDL.....	582
<i>Alexander Shleyfman, Erez Karpas</i>	
On Chatbots Exhibiting Goal-Directed Autonomy in Dynamic Environments.....	588
<i>Biplav Srivastava</i>	
Safe Goal-Directed Autonomy and the Need for Sound Abstractions.....	591
<i>Siddharth Srivastava</i>	
Exploiting Micro-Clusters to Close The Loop in Data-Mining Robots for Human Monitoring.....	595
<i>Einoshin Suzuki</i>	
Learning Abstractions by Transferring Abstract Policies to Grounded State Spaces.....	598
<i>Lawson L. S. Wong</i>	
Information-Efficient Model Identification for Tensegrity Robot Locomotion.....	602
<i>Shaojun Zhu, David Surovik, Kostas Bekris, Abdeslam Boularias</i>	

Author Index