

1. Record Nr.	UNINA9910412088803321
Autore	Conitzer Vincent
Titolo	AIES'19 : proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society : January 27-28, 2019, Honolulu, HI, USA / / Vincent Conitzer, Gillian Hadfield, Shannon Vallor
Pubbl/distr/stampa	New York, New York : , : Association for Computing Machinery, , 2019
Descrizione fisica	1 online resource (562 pages) : illustrations
Disciplina	006.3
Soggetti	Artificial intelligence Artificial intelligence - Moral and ethical aspects Artificial intelligence - Social aspects
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Sommario/riassunto	<p>It is our great pleasure to welcome you to the proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society - AIES 2019. The second edition of this conference was co-located with AAAI-19 on January 27-28, 2019 in Honolulu, Hawaii, USA. Concerns about the impact of AI on society have continued to grow in the year since AAAI and ACM joined to create the first Conference on AI, Ethics and Society. In the vision of this joint effort, it is only through multidisciplinary engagement and scholarship that we can hope to develop good responses to the challenge of ensuring that AI develops in a way that is safe and beneficial for everyone. This year's conference contributed to building a community of shared concepts and concerns. With submissions ranging across a diverse array of fields - computer science, philosophy, economics, sociology, psychology, and law - we enjoyed two days of engaging contributions that provided new paths for research. A panel and four invited talks set the stage. The opening panel on Responsible Artificial Intelligence: Aligning Technology, Engineering and Ethics (with Yolanda Gil, Huw Price, Francesca Rossi, and Dekai Wu, moderated by Peter Herschok) was held at the University of Hawaii - Manoa in the evening preceding the main conference. The</p>

main program led off with Ryan Calo's invited talk titled How We Talk About AI (And Why It Matters). He highlighted the dangers of how we talk about AI: terms like "killer robot" and "arms races" may increase rather than decrease international cooperation, and Calo argued that focusing on "ethics" may reduce our attention to formal law and governance. In her talk titled Guiding and Implementing AI, Susan Athey (Stanford) drew on the theories of causal inference and incentive design to generate insights and research paths for the design of AI systems and organizations that integrate AI and human decision makers. Anca Dragan (Berkeley) urged a reconceptualization of the challenge of building robots that achieve human objectives within constraints as a human-robot collaboration problem with the aim of both coordinating behaviour and jointly discovering the goals of collaboration, in her talk titled Specifying AI Objectives as a Human-AI Collaboration Problem. Finally, in his talk titled The Value of Trustworthy AI, David Danks (Carnegie Mellon) explored the philosophical and psychological conditions for 'trustworthy' AI systems, and the importance of this complex notion of 'trustworthiness' for informed AI policy and public action. The program of the conference included peer-reviewed paper presentations, with 35 papers presented as posters with an associated brief spotlight presentation, and another 34 as longer oral presentations. (The latter also had associated posters, if the authors so chose.) The program also included student posters. Conference session themes included algorithmic fairness, norms and explanations, artificial agency, autonomy and lethality, rights and principles, social science models for AI, measurement and justice, AI for social good, and human-machine interaction.

---