Preface

The study of consciousness remains a challenge that spans multiple disciplines. Consciousness has a demonstrated, although poorly understood, role in shaping human behavior. The processes underpinning consciousness may be crudely replicated to build better AI systems. Such a "top-down" perspective on AI readily reveals the gaps in current data-driven approaches and highlights the need for "better AI." At the same time, the process of designing AI systems creates an opportunity to better explain biological consciousness and its importance in system behavior.

Measuring the components that may lead to consciousness (e.g., modeling and assessing others' behaviors; calculating utility functions for not only an individual agent, but also an interacting society of agents) is increasingly important to address concerns about the surprising capabilities of today's AI systems.

The symposium on "Towards Conscious AI Systems" was held at Stanford University, CA, in March 25-27, 2019 as part of AAAI Spring Symposia. It was an excellent opportunity for researchers considering consciousness as a motivation for "better AI" to gather, share their recent research, discuss the fundamental scientific obstacles, and reflect on how it relates to the broader field of artificial intelligence and robotics.

Research on consciousness and its realization in AI systems motivates research to account for, with scientific rigor: the motivations of AI systems, the role of sociality with and between machines, and how to implement machine ethics.

The meeting offered a platform to discuss the connection between AI systems and other fields such as psychology, philosophy of mind, ethics, and neuroscience.

Some of the topics covered by the symposium were:

- Recent work on conscious AI systems
- Embodied conscious AI systems
- Self-reflective higher-order AI systems
- Ethical issues involving conscious AI systems
- Trust in conscious AI systems
- Social robotics and conscious AI systems
- Consciousness, the theory of mind and artificial emotions
- The role of episodic memory in conscious AI systems
- Design strategies versus developmental approaches
- Symbolic versus deep neural networks in conscious AI systems
- Measurement of consciousness in AI systems
- Physicalist models of consciousness
- Philosophy of mind and machine consciousness
- Conscious processes and time
- Implementing neuroscience of consciousness in AI systems
- Computational models of consciousness

The symposium received about 60 submission. All the submitted papers were peer reviewed by the organizing committee with the help of external reviewers. The symposium organizers would like to express their gratitude to all the colleagues for submitting papers to the symposium as well as to the colleagues that helped them in reviewing and organizing an attractive symposium.

March 2019

Antonio Chella
David Gamez
Patrick Lincoln
Riccardo Manzotti
Jonathan Pfautz

Organization

Organizing Committee

Antonio Chella	University of Palermo & ICAR-CNR, Palermo, Italy
David Gamez	Middlesex University, London, UK
Patrick Lincoln	SRI International
Riccardo Manzotti	IULM University, Milan, Italy
Jonathan Pfautz	DARPA

Additional reviewers

Igor Aleksander
Agnese Augello
Joscha Bach
Christian Balkenius
Paul Bello
Angelo Cangelosi
Cristiano Castelfranchi
Ron Chrisley
Luisa Damiano
Haris Dindo
Alessandro Di Nuovo
Salvatore Gaglio
Naveen Sundar Govindarajulu
Ignazio Infantino
David Israel
Subhash Kak
Arianna Pipitone
Roberto Pirrone
Valeria Seidita
John Sullins.