The International Association for Computing and Philosophy

2016 Annual Meeting

University of Ferrara, Department of Economics and Management
Via Voltapaletto, 7, Ferrara, Italy

June 14-17 2016

Keynote Addresses

- Aula Magna -

Mario Piazza
Associate Professor at University of Chieti-Pescara
What Does Arrow's Information Paradox Say (To Philosophers)
Tuesday June 14, 14:00-15:00

Luciano Floridi
Professor of Philosophy and Ethics of Information at University of Oxford
On Human Dignity as a Foundation for the Right to Privacy
The Copernicus Lecture
Wednesday June 15, 14:00-15:00

Jack Copeland
Professor of Philosophy at University of Canterbury
The Stored-Program Story: What Philosophy Taught History
The 2016 Covey Award and the 2016 Minds and Machines Lecture
Thursday June 16, 14:00-15:00

Marcin Milkowski
Assistant Professor at Institute of Philosophy and
Sociology of the Polish Academy of Sciences
Integrating Computational Models of the Mind
The 2016 Simon Award
Friday June 17, 11:00-12:00

Program

Tuesday June 14th
Department of Economics and Management
Via Voltapaletto, 7, Ferrara, Italy

8:30	Registration	
9:15	Opening Remarks	
	Session A - EC 6	**Session B - EC 7**
9:30	**Richard Evans**, A Kantian Cognitive Architecture	**Ugo Pagallo**, The Challenges of Digital Democracy, and How to Tackle Them in the Information Era
10:00	**Edoardo Datteri**, Model-oriented, Data-oriented, and Structural Large-scale Simulations in Neuroscience	**Anne Gerdes**, Lethal Autonomous Weapon Systems and Responsibility Gaps
10:30	**Dimitri Coelho Mollo**, Deflating Representational Content	**Don Berkich**, On the Feasibility of Autonomous Artificial Agents
11:00	Coffee Break	
11:30	**Steve Mckinlay**, Evidence, Explanation and Predictive Data Modelling	**Massimo Durante**, Security and Trust in the Digital Era
12:00	**Nicola Angius & Guglielmo Tamburrini**, From Simulation Programs as Theories to Theories of Simulation Programs	**Johnny Hartz Søraker**, A Right to be Remembered? Meaningfulness and Permanence in Virtual Environment
12:30	**Russ Abbott**, Reductionism vs Reverse Engineering	**Erica Neely**, The Ethics of Choice in Single-Player Video Games
13:00	Lunch	
14:00	Keynote Address: **Mario Piazza**	What Does Arrow's Information Paradox Say (To Philosophers)
15:00	Coffee Break	
15:30	**Concurrent Symposia**	
EC 6	**Gualtiero Piccinini**, Computation and Representation in Cognitive Neuroscience	Gualtiero Piccinini, Carl Craver, Andrea Scarantino, Nir Fresco, David Kaplan, and Oron Shagrir
EC 7	**David Danks & Heather Roff**, Automated and Autonomous Conflicts: AI, Ethics, and the Conduct of Hostilities	David Danks, Heather Roff, Ugo Pagallo, and Mariarosaria Taddeo
18:30	End of the Sessions	
19:00	Wine Reception: Bar Schifanoia, Palazzo Schifanoia, Via Scandiana, 21	

Program

Wednesday June 15th
Department of Economics and Management
Via Voltapaletto, 7, Ferrara, Italy

8:30	Registration	
	Session A - EC 6	**Session B - EC 7**
9:30	**Orlin Vakarelov**, Information Media Networks: A Representational Framework for Cognition	**Tony Doyle**, Obfuscation and Strict Online Anonymity
10:00	**Hasen Khudairi**, Modal Ω-Logic: Automata, Neo-logicism, and Set-theoretic Realism	**Alexis Elder**, Taking Control of the Conversation: Autonomy and Computer-Mediated Communication
10:30	**Marcello D'Agostino**, Logic, Semantic Information, and Computational Complexity	**Michele Rapoport**, Dwelling Technologically: Territory and Labor in the Automated Home
11:00	Coffee Break	
11:30	**Paul Schweizer**, Computation in Physical Systems: a Normative Mapping Account	**Frances Grodzinsky, Marty J. Wolf, and Keith Miller**, Applying a Social-Relational Model to Explore the Curious Case of hitchBOT
12:00	**Vincenzo Fano, Pierluigi Graziani, Mirko Tagliaferri, and Gino Tarozzi**, When is a Computation Realized by a Physical System?	**Gordana Dodig Crnkovic & Gaetana Sapienza**, Ethical Aspects of Technology in the Multi-Criteria Decision Analysis
12:30	**Mark Addis & Doukas Kapantaïs**, The Church Turing Thesis and Computability in Peano Arithmetic	**Owen King**, Integrating Responsibility into the Training Data for Machine Learning Systems
13:00	Lunch	
14:00	Keynote Address: **Luciano Floridi**	On Human Dignity as a Foundation for the Right to Privacy
15:00	Coffee Break	
15:30	**Concurrent Symposia**	
Aula Magna	**Selmer Bringsjord & John Licato**, New-Millennium Logic, Computing, & the Mind	Selmer Bringsjord, John Licato, Paul Bello, Matthias Scheutz, Kevin O'Neill
EC 6	**Sabine Thürmel & Judith Simon**, Ethics of Big Data: The Engineering of the "Not Yet"	Wolfgang Pietsch, Gernot Rieder, Brent Mittelstadt, and Sabine Thürmel
EC 7	**Nicola Angius, Giuseppe Primiero, and Petros Stefaneas**, Methodological Issues in the Philosophy of Computer Science	Raymond Turner, Petros Stefaneas, Robin Hill, Giuseppe Primiero, Viola Schiaffonati, Mario Verdicchio, and Nir Fresco
18:30	End of the Day	

Program

Thursday June 16th
Department of Economics and Management
Via Voltapaletto, 7, Ferrara, Italy

8:30	Registration	
	Session A - EC 6	**Session B - EC 7**
9:30	**Hyungrae Noh**, Shannon, Dretske and Millikan on Information and Communication	**Jason Borenstein & Ronald Arkin**, Robots, Ethics, and Intimacy: The Need for Scientific Research
10:00	**Matteo d'Alfonso**, Virtual information in the light of German Idealism	**Anna Wilks**, Robotic Responsibility
10:30	**Francois Oberholzer & Stefan Gruner**, The Notion of 'Information': Enlightening or Forming?	**Ioan Muntean & Don Howard**, Generalized Moral Agency: Artificial Moral Functionalism and Moral Behaviorism
11:00	Coffee Break	
11:30	**Patrick Allo**, Modelling Online Collaborative Mathematics	**Migle Laukyte**, Against Human Exceptionalism: Environmental Ethics and Machine Question
12:00	**Sandra Wachter**, Privacy – What is it good for? Absolutely nothing?	**Teresa Numerico**, Big Data
12:30	**Daniele Porello**, Judgment Aggregation in Relevant Logic	**Thomas M. Powers & Jean-Gabriel Ganascia**, Towards a Prescriptive Deontological Non-monotonic Machine Ethics
13:00	Lunch	
14:00	Keynote Address: **Jack Copeland**	The Stored-Program Story: What Philosophy Taught History
15:00	Coffee Break	
15:30	**Concurrent Symposia**	
EC 6	**Don Berkich**, Lightning Rounds	
EC 7	**Gordana Dodig-Crnkovic**, Embodied Cognition: Constructivist and Computationalist Perspectives	Robert Lowe, Gordana Dodig-Crnkovic, Alexander Almér, Rickard von Haugwitz, Ron Chrisley, Tom Ziemke, Vincent Müller, Piotr Bołtuć, Marcin Schröder, Marcin Milkowski, and Oron Shagrir
19:00	End of the Sessions	
20:00	Banquet: Ristorante Big Night, Via Largo Castello, 38	

Program

Friday June 17th
Department of Economics and Management
Via Voltapaletto, 7, Ferrara, Italy

8:30	Registration	
	Session A - EC 6	**Session B - EC 7**
9:30	**Luca Rivelli**, Antimodularity: Pragmatic Consequences of Computational Complexity on Scientific Explanation	**Richard Heersmink**, Does the Web Diminish or Enhance Our Cognitive Skills?
10:00	**Ingvar Tjostheim**, Telepresence and the Role of the Senses	**James Williams**, Three Types of Distraction
10:30	Coffee Break	
11:00	Keynote Address: **Marcin Milkowski**	Integrating Computational Models of the Mind
12:00	Closing Remarks	
12:30	IACAP Business and Planning Meeting	
13:00	End of the Conference	

Keynote Addresses

(in order of presentation)

Mario Piazza

Associate Professor at University of Chieti-Pescara
"What Does Arrow's Information Paradox Say (To Philosophers)"
Tuesday June 14, 14:00-15:00

Kenneth Arrow's information paradox (AIP, henceforth) diagnoses an inherent conflict between two parties inclined to exchange information: in order for the prospective buyer to ascertain the value of the information, she needs to know its content; but once the seller discloses the information, the buyer has no reason to buy it any longer. While practical implications of AIP have been widely discussed in patent and economic literature, it seems to me that AIP also accomplishes real but underexplored theoretical work about the nature of information. This talk intends to illustrate some aspects of the theoretical significance of AIP. Moreover, I discuss the intriguing conceptual relation between AIP and the notion of zero-knowledge proofs in cryptography: roughly speaking, these are protocols that enable a prover to convince a verifier that a statement is true, without conveying any additional information. Some moral relevant to philosophy, and in particular to philosophy of information, is drawn. ↑**Program**↑

Luciano Floridi

Professor of Philosophy and Ethics of Information at University of Oxford
"On Human Dignity as a Foundation for the Right to Privacy"
The Copernicus Lecture
Wednesday June 15, 14:00-15:00

Human dignity is the fundamental framework within which one needs to interpret what the European General Data Protection Regulation – and more generally European culture and jurisdiction – understand by informational privacy. The protection of privacy should be based directly on the protection of human dignity, not indirectly, through other rights such as that to property or to freedom of expression. However, unless one explains convincingly what human dignity may mean in the twenty-first century and in a mature information society, it remains obscure and questionable exactly which interpretation of human dignity may provide the foundation for privacy (as well as all other human rights), and hence why. This is the challenge that I address in this talk. ↑**Program**↑

Jack Copeland

Professor of Philosophy at University of Canterbury
"The Stored-Program Story: What Philosophy Taught History"
The 2016 Covey Award and the 2016 Minds and Machines Lecture[1]
Thursday June 16, 14:00-15:00

This lecture describes the early history of a fundamental ingredient of modern computing, the stored-program concept. An analysis of the contributions by Alan Turing, Konrad Zuse, F. C. Williams, Tom Kilburn, Presper Eckert, John Mauchly, Richard Clippinger and John von Neumann shows that the stored-program concept consisted of several distinct layers or 'onion-skins'. These different onion-skins emerged slowly over a ten-year period, giving rise to a number of different programming paradigms. The 'onion-skin analysis' permits the resolution of various scholarly disagreements about the stored-program concept and its history. ↑**Program**↑

[1]Sponsored by Springer International Publishing

Marcin Milkowski
Assistant Professor at Institute of Philosophy and Sociology of the Polish Academy of Sciences
"Integrating Computational Models of the Mind"
The 2016 Simon Award
Friday June 17th, 11:00-12:00

A vast majority of theoretical papers in cognitive science today describe computational models of cognitive processes. My focus in this talk will be on attempts to integrate separate computational models of the mind. Most models describe just how human and non-human subjects solve particular cognitive tasks. Since the 1970s, modelers are, however, acutely aware of the fact that partial models of particular tasks do not simply add up. There are several strategies that are supposed to help integrate them, or fill the gaps between these partial explanations. I will analyze three: cognitive architectures, interfaces, and experimental constraints.

The classical strategy of integration, pursued by proponents of production systems, is to describe complete cognitive architectures reused in multiple explanations. In this case, theorists usually present simplified versions of previously offered explanations, by showing how they can be integrated in their architectures. For example, it's possible to use Baddeley's account to model working memory in ACT-R. In the process, modelers implement features that were not in the original theory, and show implications for its further development. In other words, previous work is reused and adapted in integrated models of cognition.

The second, contemporary strategy is to build interfaces between different computational models and systems. For example, one may connect NEURON simulation software with LFPy, Python simulation package that computes biophysical properties of Local Field Potentials. In such a case, the user simply can use LFPy to compute values necessary for his or her simulation of neural networks. A much more interesting case is building theoretically inspired interfaces between various computational subsystems, for example in hybrid cognitive architectures.

The last, much weaker, strategy is to constrain cognitive models by respecting certain known psychological limitations. For example, Herbert Simon and Allen Newell presupposed that the system may not have working memory bigger than 7 ± 2 meaningful chunks. It's also presupposed that the model may not execute more than 100 operations per second, which is estimated to be the computational speed of the nervous system.

I will consider the advantages and disadvantages of these strategies, in particular how they can contribute to building integrated mechanistic models of cognition. ↑**Program**↑

Concurrent Symposia

(in order of presentation)

Gualtiero Piccinini

"Computation and Representation in Cognitive Neuroscience"
Gualtiero Piccinini, Carl Craver, Andrea Scarantino, Nir Fresco, David Kaplan, and Oron Shagrir
Tuesday June 14, 15:30-18:30, EC 6

Cognitive neuroscientists routinely explain cognition in terms of neural computations over neural representations. Yet the notions of neural representation and neural computation remain poorly understood. There is no consensus on how to construed these notions and how to relate them to the notions of computation and representation used in other disciplines, including psychology and computer science. In this symposium, some of the leading philosophers who work on foundations of cognitive neuroscience will come together to push the debate forward on these central concepts. ↑**Program**↑

David Danks & Heather Roff

"Automated and Autonomous Conflicts: AI, Ethics, and the Conduct of Hostilities"
David Danks, Heather Roff, Ugo Pagallo, and Mariarosaria Taddeo
Tuesday June 14, 15:30-18:30, EC 7

Conflicts—whether overt warfare or hostile actions that fall short of that threshold—are increasingly automated, and increasingly governed by autonomous systems. In the physical domain, autonomous weapons systems (including sophisticated sensor fusion), automated war gaming, and battle management software all involve transfer of (some) control to machines. In the electronic domain, information warfare, cyber-hostilities, and mass surveillance all require some degree of automation, whether because of the speed, complexity, or scale of the conflicts.

There have been numerous debates in academia, government, and the public sphere over the value and desirability of increasingly autonomous technologies; much of this debate has focused on their moral or legal permissibility. Surprisingly, however, these debates have rarely examined the ways in which artificial intelligence can fundamentally change the conduct of warfare and other conflicts. AI techniques and autonomous systems are not simply additional pieces on a fixed board; instead, they arguably change the very rules, norms, and practices of conflict. As a result, those systems can give rise to distinctive, novel moral problems for future battlefields and conflict spaces.

This symposium brings together an interdisciplinary set of scholars to look at the novel moral challenges, and broader moral implications, of developing and deploying increasingly automated and autonomous systems that rely on artificial intelligence. This symposium will focus on conflicts in both the physical and electronic domains, as well as both warfare and hostilities short of warfare. We expect symposium participants to examine not only weapons systems, but also cybersecurity and defense, intelligence surveillance and reconnaissance, targeting, and command & control. This symposium will touch on topics ranging from ethical applications of AI, to the changing face and scope of warfare, to the challenges of automation and autonomous systems in conflict situations. Talks in the symposium will cover multiple core topics and areas of the IACAP-16 conference. As such, we expect that it will be of wide interest and relevance for multiple attendees. ↑**Program**↑

Selmer Bringsjord & John Licato

"New-Millennium Logic, Computing, & the Mind"
Selmer Bringsjord, John Licato, Paul Bello, Matthias Scheutz, Kevin O'Neill
Wednesday June 15, 15:30-18:30, Aula Magna

Aristotle inaugurated formal deductive logic before, and for exploration in, the first millennium. The Organon provided principally two things: an initial (and painfully inadequate) stab at formalizing the stunningly seminal reasoning of Euclid, and another at formalizing the capacity for humans to acquire and reason deductively over inferentially interconnected declarative knowledge and belief.

The first of these two achievements eventuated, near the end of the second millennium, in perhaps modern mathematical logic's greatest triumph: a formalization of not only Euclid's reasoning, but the reasoning of mathematicians through the ages. In addition, since this formalization allowed the Entscheidungsproblem to be rigorously posed, and settled, general-purpose computation at the level of standard Turing machines (and their equivalents) arrived on the scene; soon these purely abstract machines were physicalized. These remarkable developments eclipsed the invention, progress in, and great promise of inductive logic, in the second millennium.

The second trajectory established by Aristotle eventuated, also near the end of the second millennium, in the modern, classical conception of AI: viz., AI as the attempt to engineer artificial intelligence by building a machine able to itself use logic at the human level. Such machines were often, and sometimes still are, called "knowledge-based" systems.

In addition, traditionally, formal logic and logicist disciplines such as decision theory have been devoted to confronting, and at least attempting to resolve, paradoxes. Aristotle, in this regard, was active. He for instance took on Zeno's paradox of the arrow. His proposed solution, while seminal, was ultimately inadequate — but when Leibniz and Newton invented the differential and integral calculus in the second millennium, the paradox of the arrow was put to rest. It's reasonable to say the same sanguine thing about other deductive paradoxes, for example The Liar. Yet many paradoxes remain unsolved to this day; particularly challenging ones produce their counter-intuitive results on the strength not just of deduction, but of induction as well, and they often involve propositional attitudes like knowledge and belief.

We are now firmly into the third millennium. Today, AI isn't just an interesting thread within the march forward of science and engineering; no, AI seems to dominate the headlines. But there are some noteworthy twists: One, robots, only creatures of fiction for the vast majority of the second millennium, are here, and fast becoming ubiquitous. Two, that part of logic-less AI based in statistical learning is all the rage: "deep learning," if media reports of today were to be trusted, will be bringing us computing machines of superhuman intelligence in just a few short years. ↑**Program**↑

Sabine Thürmel & Judith Simon

"Ethics of Big Data: The Engineering of the "Not Yet""
Wolfgang Pietsch, Gernot Rieder, Brent Mittelstadt, and Sabine Thürmel
Wednesday June 15, 15:30-18:30, EC 6

Big data approaches are deployed in many sectors of contemporary life, such as the military, civil surveillance, the online economy, the workplace or the education and health systems. In public environments the engineering of the "not yet" is for instance employed to improve the public infrastructure using "smart city approaches" or to prevent potentially undesirable developments in the case of "predictive policing". While number crunching and statistical calculations form the basis of big data practices, it is in particular the temporality of data practices that is one of the most interesting features of big data and of central interest to

our symposium. Big data practices aim at generating hypotheses about the future based on past or present data. However, these hypotheses do not merely provide predictions about the future, they also have an impact on the future when used to engineer the "not yet". In our interdisciplinary symposium, we aim to address this temporality of big data practices from multiple perspectives. The participants will give short kick-off presentations followed by a roundtable discussion. ↑**Program**↑

Nicola Angius, Giuseppe Primiero, and Petros Stefaneas

"Methodological Issues in the Philosophy of Computer Science"
Raymond Turner, Petros Stefaneas, Robin Hill, Giuseppe Primiero, Viola Schiaffonati, Mario Verdicchio, and Nir Fresco
Wednesday June 15, 15:30-18:30, EC 7

In the last decade, the Philosophy of Computer Science (PhiloCS) arose with the aim of developing essential epistemological research in computer science and software engineering, providing an analysis of methods and techniques involved in the software development processes. This symposium aims at presenting some of the current research trends in PhiloCS and to show the impact that a methodological examination of computer science may have, in particular, on the philosophy of cognitive sciences and AI. The symposium will deal with the following three main topics: specifications and miscomputations of computational artefacts; design, representation and testing of computational artefacts; the impact of PhiloCS on computability, cognitive sciences and AI. ↑**Program**↑

Don Berkich

"Lightning Rounds"
Thursday June 16, 15:30-19:00, EC 6

At IACAP's 2014 meeting in Thessaloniki I introduced a symposium, "The Lightning Rounds", consisting of 5-minute presentations on any topic of the presenter's choice, followed by three minutes question and answer and a two minute break for discussion and set-up before moving on to the next Lightning Round. I solicited presenters from among conference attendees the day prior to the session. Popular at computer science conferences, I described the symposium as an experiment to see whether such an event would succeed at an IACAP meeting. Attendance varied from 5 to 12 during the two-hour event, with a total of 8 presentations. Presentations were lively, witty, varied, and fascinating. Discussion was vigorous. Upon informal polling, attendees and presenters alike unanimously recommended the Lightning Rounds be made a regular feature.

Accordingly, I proposed and received approval to further develop the Lightning Rounds at IACAP's 2015 meeting in Delaware. The Lightning Rounds in Delaware were, I thought, highly successful. Attendance varied from 14 to 18, and we had enough presentations to completely fill 2 hours with a total of 10 presentations. Situating the Rounds at the end of the conference worked extremely well, as presenters had much to say about what they had heard and discussed the previous few days at the conference. Moreover, I'm increasingly convinced that the short format presentations are a refreshing change of pace for audience members and a boon to presenters in terms of trying out arguments and ideas they'd otherwise refrain from presenting.

As before, it was an engaging, fast-paced, and thoroughly enjoyable 'cap' to the conference. My plan in Ferrara is to solicit six or eight presentations in advance and leave the remaining slots open for participants. Ideally, the Lightning Rounds would again be timed well towards the end of the conference to encourage commentary and counter-argument on conference presentations. ↑**Program**↑

Gordana Dodig-Crnkovic

"Embodied Cognition: Constructivist and Computationalist Perspectives"
Robert Lowe, Gordana Dodig-Crnkovic, Alexander Almér, Rickard von Haugwitz, Ron Chrisley, Tom Ziemke, Vincent Müller, Piotr Bołtuć, Marcin Schröder, Marcin Milkowski, and Oron Shagrir
Thursday June 16, 15:30-19:00, EC 7

The term cognition is a controversial one when considered in terms of its constitution and function. On the one hand, it has been construed as a type of computation consisting of syntactical manipulations of symbolic representations. At the opposite extreme cognition is considered to be continuous with life; its constitution being constantly re-established by autopoietic self-organization fundamental to all living systems. Many perspectives, however, reside somewhere between these extremes – advocates of constructivist (including enactivist) approaches are not universal in their eschewing of representational language, and computationalists are not only those who construe cognition as "language of thought" or even human language processing.

An aspect of cognition central to the above divide is Embodiment. Embodied cognition holds that cognition is grounded in environmental interactions in the world and is invisible in classical symbolic representation accounts of cognitive function, which is modeled on human "thinking" or "mentality". However, modern computational perspectives on cognition such as natural computation (including info-computationalism) account for embodiment whereby cognitive processes are considered to emerge from interactions in the world. In this symposium, we wish to encourage frank debate about the perceived differences in the various perspectives on constructivist and computationalist accounts of cognition, and specifically embodied cognition. This will be fostered by the balanced representation of speakers and panel discussants that represent (often diametrically) opposing perspectives in the area. This debate concerns the Cognitive Science, Computation & Cognition theme of the IACAP annual meeting and provides critical arguments concerning the controversies regarding the nature of cognition. Could it be that different approaches focus on different aspects of cognition? Thinking in early computationalism vs. generative and evolutionary mechanisms in embodied cognition? Is it possible to reconcile constructivism with computationalism in a new synthesis? What is the role of emotions in computational approaches? What is the role of higher cognitive functions in embodied approaches? ↑**Program**↑

Papers (in alphabetical order by first author)

Russ Abbott

California State University-Los Angeles
"Reductionism vs Reverse Engineering"
Tuesday June 14th, 12:30, Session A

Albert Einstein and Steven Weinberg claim that all of science can be deduced from the fundamental laws of physics. Philip Anderson and Erwin Schrödinger disagree saying that nature is constructive in ways that go beyond reductionism. Both sides may be right: physics may entail the inevitability of nature's regularities—but not necessarily their precise form. We suggest that science may be best understood as constructive analysis or the reverse engineering of nature. ↑**Program**↑

Mark Addis & Doukas Kapantaïs

Birmingham City University, Academy of Athens
"The Church Turing Thesis and Computability in Peano Arithmetic"
Wednesday June 15th, 12:30, Session A

An interpretation of the Church Turing thesis as a conjecture which claims anything that can be effectively calculated can be computed by a Turing machine equivalent with all maximal models of computation being equivalent to Turing machines up to isomorphism is considered. A method for showing that an idealised human calculator can prove the totality of the Ackermann function (which is based on a proof that the function terminates for all arguments) whilst no Turing machine or its equivalent can do this is provided. The algorithmic method for this computation is neither Turing machine equivalent reducible nor describable but it is intuitively clear that it is both effective and mechanical in nature. Related arguments to these about the totality of the Ackermann function being human calculator but not Turing machine computable are given for Goodstein's theorem. These arguments about the totality of the Ackermann function and Goodstein's theorem refute the conjecture interpretation of the Church Turing thesis that anything which can be effectively calculated can be computed by a Turing machine equivalent with all maximal models of computation being equivalent to Turing machines up to isomorphism. This conclusion raises doubts about whether the Church Turing thesis in general has much value since only the isomorphism interpretation is informative about the nature and structure of computability. ↑**Program**↑

Patrick Allo

University of Oxford
"Modelling Online Collaborative Mathematics"
Thursday June 16th, 11:30, Session A

In this paper insights from formal methods that can be used to model scientific communities are used to complement insights from a computational analysis and the application of some elements from social network theory to a prominent case of mathematical online collaboration, namely the Polymath Projects. ↑**Program**↑

Nicola Angius & Guglielmo Tamburrini
Università di Sassari & Università di Napoli
"From Simulation Programs as Theories to Theories of Simulation Programs"
Tuesday June 14th, 12:00, Session A

This paper provides a methodological analysis of Executable Cell Biology (ECB) applying formal verification methods developed in theoretical computer science to current computational biology simulative inquiries. First, it is highlighted how ECB resumed the general idea of constructing theoretical models that are also executable, pursued over fifty years ago by Allen Newell and Herbert Simon in the specific IPP methodology. However, instead of a simulative program, ECB focuses on a more abstract model of the biological system. On one hand, the processes of abstraction involved in the construction of ECB theoretical models allow to omit those implementation details of the simulative program that have no theoretical value. On the other hand, the executability of the abstract model permits to expand the class of predictions that can be extracted from the observation of the simulative programs' executions. Secondly, ECB executable theoretical models which are distinct from the simulative programs are here acknowledged as mechanism sketches providing functional explanations of simulated cell systems' behaviours. It is shown how distortive idealization processes involved in the construction of those functional models prevent one from making correct functional role instantiations and consequently from developing full-fledged mechanist explanations. ↑**Program**↑

Don Berkich
Texas A&M University-Corpus Christi
"On the Feasibility of Autonomous Artificial Agents"
Tuesday June 14th, 10:30, Session B

Our inability to describe a computational solution to the hard problem of phenomenal consciousness does not preclude the existence of a solution, but it does curtail the aspirations one might entertain for engineering minds. How much else is at stake, though? In this paper I explore the extent to which the apparent hard-limit of phenomenal consciousness likewise curtails computational accounts of agency. Framed in terms more familiar to the philosophy of mind, I conclude that while philosophical zombies can be agents, they cannot be autonomous agents. The hard problem of phenomenal consciousness is thus a vicious computational contagion: Autonomous agency, moral responsibility, and personhood all join phenomenal consciousness in the quarantine of the computationally intractable. ↑**Program**↑

Jason Borenstein & Ronald Arkin
Georgia Institute of Technology
"Robots, Ethics, and Intimacy: The Need for Scientific Research"
Thursday June 16th, 9:30, Session B

Intimate relationships between robots and human beings will begin to form in the near future. Market forces, customer demand, and other factors will drive the creation of various forms of robots to which humans may form strong emotional attachments. Yet prior to the technology becoming fully actualized, numerous ethical, legal, and social issues must be addressed. This could be accomplished in part by establishing a rigorous scientific research agenda in the realm of intimate robotics, the aim of which would be to explore what effects the technology may have on users and on society more generally. Our goal is not to resolve whether the development of intimate robots is ethically appropriate. Rather, we contend that if such robots are going to be designed, then an obligation emerges to prevent harm that the technology could cause. ↑**Program**↑

Marcello D'Agostino
University of Milan
"Logic, Semantic Information, and Computational Complexity"
Wednesday June 15th, 10:30, Session A

We discuss the impact of the theory of computational complexity on the received view about the relationship between logic and semantic information. Next, we outline a more realistic approach drawing upon a kind of informational semantics for the logical operators that was essentially anticipated by W.V.O. Quine, but ignored by the logical community. This naturally leads to an infinite hierarchy of tractable consequence relations, defined in terms of "depth-bounded" information states, that converge to classical propositional logic. We argue that this approach provides an informational view of classical logic that allows for modeling the deductive components of non-ideal reasoning agents. ↑**Program**↑

Matteo d'Alfonso
Università di Ferrara
"Virtual information in the light of German Idealism"
Thursday June 16th, 10:00, Session A

In (D'Agostino, Floridi 2009) the authors face the so-called "scandal of deduction" (Hintikka, 1973). This lies in the fact that Bar-Hillel-Carnap theory of semantic information implies that tautologies carry no information. Given that any mathematical demonstration and more in general every logical inference in a first-order language can be reduced to a tautology; this would imply that demonstrations bring no fresh information at all. Addressing this scandal (D'Agostino, Floridi 2009) offers both: i) a logic model for a strictly analytical reasoning, where the conclusions depend just on the information explicitly present in the premises; and ii) a proposal for the ranking of the informativeness of deductions according to their increasing recourse to so called "virtual information", namely information that is temporarily assumed but not contained in the premises. In this paper I will focus on the status of virtual information in its connection with the Kantian philosophical spirit. Exploiting the standard Kantian difference between theoretical and practical reason, my aim is to show that the access to virtual information is due to what Kant calls practical reason rather then to the theoretical one, even though the effects of its deployment are purely theoretical, i.e. don't lead an agent to any moral action but just to acquiring new information. ↑**Program**↑

Edoardo Datteri
Università degli Studi di Milano-Bicocca
"Model-oriented, Data-oriented, and Structural Large-scale Simulations in Neuroscience"
Tuesday 14th, 10:00, Session A

In this short paper I argue that contemporary large-scale brain simulation projects illustrate various roles played by simulations in the study of the brain, over and above the "synthetic method" widely used in cognitive science and biorobotics. I distinguish between behavioral and structural simulations, and between model-oriented and data-oriented simulation studies. Moreover, I argue that these distinctions can be usefully brought to bear on the question whether high biological accuracy is to be praised in a large-scale simulation study or not. Overall, I believe that the proposed analysis – to be refined in future research – can be useful to assess the epistemic value of large-scale computer and robotic simulations in the study of the brain and in the explanation of behavior. ↑**Program**↑

Gordana Dodig Crnkovic & Gaetana Sapienza

Mälardalen University, ABB Corporate Research
"Ethical Aspects of Technology in the Multi-Criteria Decision Analysis"
Wednesday June 15th, 12:00, Session B

In technological systems, decisions are often governed by multi criteria decision analysis (MCDA) techniques that take into account mutually opposing criteria for the system, and it results in ranking of alternatives. MCDA is based on value systems of decision-makers, and ethical deliberation in the process is implicit. We argue that it is necessary to make decision-making in technological systems transparent such that value basis and ethical considerations become explicit and subject for scrutiny of involved stakeholders. As different priorities, value systems and ethical choices result in different technical solutions, such solutions when put in use will promote those intrinsic and implicit values. In a society with ubiquitous technology, value aspects of technology are essential. At present there is no explicit mechanism to expose ethical aspects in these analyses, so they can easily be forgotten. As a support to encourage introduction of transparent value-based deliberation we propose an extended MCDA scheme that explicitly takes into account ethical analysis. ↑**Program**↑

Tony Doyle

Hunter College
"Obfuscation and Strict Online Anonymity"
Wednesday June 15th, 9:30, Session B

As our digital wake ripples out, big data is standing by to ride it. The mass collection, aggregation, analysis, and dissemination of personal information permit unnerving inferences about our characters, preferences, and future behavior that were inconceivable just a couple of decades ago. This paper looks primarily at online searching and the commercial harvesting of personal information there. I argue that our best hope for protecting privacy online is anonymity through obfuscation. Obfuscation attempts to throw data collectors off one's digital trail by making personal data less useful. However, anonymous web searching has costs. I examine two of the most serious and urge that they are worth paying in the light of the heavy toll the commercial gathering and analysis of our information takes on privacy and autonomy. I close with some thoughts on (1) how individual, rational decisions have led to a surveillance regime that few would have chosen beforehand and (2) the alleged autonomy of information technology. ↑**Program**↑

Massimo Durante

University of Turin
"Security and Trust in the Digital Era"
Tuesday June 14th, 11:30, Session B

Security is a fundamental lever that allows a government to gain greater powers of direction and control with the aim of immunizing society against the risks that threaten its integrity. The author argues that security is neither managed nor granted on one's own, but is rather delegated through fiduciary forms that institutionalise the risk-management. Security, therefore, needs trust. However, where trust is, there is risk. Today, the risk against which society claims to immunized emerges again with the new, increasingly mediated by Technologies and less and less politically legitimized, forms of fiduciary management, where the possibility of a weakening of rights and of political responsibility resides. ↑**Program**↑

Alexis Elder

Southern Connecticut State University
"Taking Control of the Conversation: Autonomy and Computer-Mediated Communication"
Wednesday June 15th, 10:00, Session B

In this paper, I investigate the impact boundary-promoting communication technology – such as texts, comments, microblogging, and instant messaging – have on friendships, and arrive at the surprising conclusion that these technologies are, despite appearances, good for personal relationships, and thereby good for us. ↑**Program**↑

Richard Evans

Imperial College London
"A Kantian Cognitive Architecture"
Tuesday June 14th, 9:30, Session A

This paper describes a computer implementation of a cognitive architecture inspired by Kant's First Critique. The computer model has been applied to a verbal reasoning task: Hofstadter's "Seek Whence" problem-set. It achieves human-level performance in this domain, significantly out-performing previous (non-Kantian) approaches. In this paper, I outline the computer model, and describe the underlying Kantian cognitive architecture. ↑**Program**↑

Vincenzo Fano, Pierluigi Graziani, Mirko Tagliaferri, and Gino Tarozzi

University of Urbino, University of Chieti-Pescara, University of Urbino, University of Urbino
"When is a Computation Realized by a Physical System?"
Wednesday June 15th, 12:00, Session A

The aim of this paper is to address the question: when is a computation realized by a physical system? Our strategy is to propose a new definition of realization that makes the above question more tractable and easier to scrutinize. We then show that our definition has some advantages when dealing with classical arguments moved against definitions of this type. The paper will be structured in four parts: in part one, we will set some desiderata our definition ought to fulfil; in part two, we will introduce some prerequisite notions that are necessary for a full comprehension of our definition; in part three, we will introduce our definition of realization – it will turn out that our definition can be identified by a specific kind of strategy that Piccinini dubs nomological mapping account; in part four, we will assess our definition. ↑**Program**↑

Anne Gerdes

University of Southern Denmark
"Lethal Autonomous Weapon Systems and Responsibility Gaps"
Tuesday June 14th, 10:00, Session B

Weaponized autonomous technologies are entering the scene of warfare apace, and all though the capabilities of warfare systems already by now allow for autonomous actions within restricted contexts, we have not yet witnessed these systems initiating lethal action. But this may very well change in a near future. This paper sets out to clarify issues of responsibility in human-technology relations in general (sec. 2), as a springboard for a discussion of the implications of responsibility gaps in relation to weaponized autonomous technologies (sec.3). Finally, I conclude that a warfare context requires a strong interpretation of responsibility, incompatible with the acceptance of responsibility gaps. Therefore, I argue that we ought not to delegate lethal decisions to autonomous weapon systems. ↑**Program**↑

Frances Grodzinsky, Marty J. Wolf, and Keith Miller
Sacred Heart University, Bemidji State University, University of Missouri-St. Louis
"Applying a Social-Relational Model to Explore the Curious Case of hitchBOT"
Wednesday June 15th, 11:30, Session B

We use the recent example of hitchBOT, a robot designed to rely on humans for transportation as it hitchhiked across Europe and North America. We explore Mark Coeckelbergh's social-relational model as it applies to hitchBOT-like robots, robots that people encounter in a casual way when no human operator is present. We will analyze issues raised by this model not only from interactions between the public and robots, but also from the point of view of developers of the interface of the robot, whose choices impact how people experience a robot and the moral status they afford it. In addition we draw on empirical work by Kate Darling on framing, that is how the attitudes of those who encounter a robot are influenced by stories about it. We show how different sorts of robots lead to different sorts of social relationships and that these differences are ethically important. As a particularly interesting example, we consider the relationships that hitchBOT's developers wished to encourage. While Coeckelbergh's and Darling's works are helpful in analyzing human/robot ethical issues, we argue that their insights do not remove all potential difficulties with those relationships or with our ability to clearly understand their ethical importance. We argue that a focus on the social relationships between humans and robots does not remove responsibility from a robot developer or owner for that robot's behavior. On the contrary, knowing that humans may form significant relationships with these machines increases developers' ethical responsibilities to ensure (as much as is practical) that those relationships will be positive, especially when there is potential for the robot to encounter other people when its owner is not present. Finally, we note that hackers greatly complicate both the situation of human/robot relationships and the resulting analysis of the ethical ramifications of these relationships. ↑**Program**↑

Richard Heersmink
Macquarie University
"Does the Web Diminish or Enhance Our Cognitive Skills?"
Friday June 17th, 9:30, Session B

We increasingly use the Web as an external memory system, in that way transforming our memory strategies. Research suggests that when we know information is accessible in some external media, we tend to put less effort into encoding it internally. Instead, we develop ways to remember where the information is stored rather than its exact content. Some argue that this strategy is adaptive because it frees up internal memory which can then be used for other (perhaps more important) cognitive tasks, whereas others argue that this is maladaptive because it makes us less knowledgeable. In this paper, I first have a closer look at the empirical evidence and suggest that we should not jump to conclusions. I then place it in a broader cognitive-historical framework and argue that in our recent evolution we have always relied on external information to perform some of our cognitive tasks, augmenting and shaping our cognitive abilities. We are thus essentially a soft self, coevolving with the tools and technologies we use. Our goal should not be to merely point the finger at the technology, but to actively and responsibly shape this co-evolutionary process in a way suitable for cognitive agents in an information society. ↑**Program**↑

Hasen Khudairi
Arché Philosophical Research Centre, University of St Andrews
"Modal Ω-Logic: Automata, Neo-logicism, and Set-theoretic Realism"
Wednesday June 15th, 10:00, Session A

This essay examines the philosophical significance of Ω-logic in Zermelo-Fraenkel set theory with choice (ZFC). The dual isomorphism between algebra and coalgebra permits Boolean-valued algebraic models of ZFC to be interpreted as coalgebras. The modal profile of Ω-logical validity can then be countenanced within a coalgebraic logic, and Ω-logical validity can be defined via deterministic automata. I argue that the philosophical significance of the foregoing is two-fold. First, because the epistemic and modal profiles of Ω-logical validity correspond to those of second-order logical consequence, Ω-logical validity is genuinely logical, and thus vindicates a neo-logicist conception of mathematical truth in the set-theoretic multiverse. Second, the foregoing provides a modal-computational account of the interpretation of mathematical vocabulary, adducing in favor of a realist conception of the cumulative hierarchy of sets. ↑**Program**↑

Owen King
Oberlin College
"Integrating Responsibility into the Training Data for Machine Learning Systems"
Wednesday June 15th, 12:30, Session B

Just as humans can make judgments responsibly or irresponsibly, so too can computers. Machine learning systems that are trained on data sets that include irresponsible judgments are likely to produce irresponsible judgments as outputs. In this paper I focus on a particular kind of judgment a computer system might make: identification of what action a human agent has performed. And I focus on a specific sort of evidence that might be the basis for such a judgment: photographic images. I identify a particular sort of irresponsibility such judgments might manifest, viz., a certain kind of presumptuousness. I explore the possibility that carefully procuring the training data for image recognition systems could ensure that the systems avoid presumptuousness. However, I observe that this sort of approach potentially retards the development of these sorts of machine learning systems in undesirable ways. The more general point of this paper extends beyond just image recognition systems and the challenge of responsibly identifying actions. The paper is an attempt to begin worrying about the ethical consequences, not just of the operation of machine learning systems, but also of the particular streams of data on which they are trained. ↑**Program**↑

Migle Laukyte
Universidad Carlos III de Madrid
"Against Human Exceptionalism: Environmental Ethics and Machine Question"
Thursday June 16th, 11:30, Session B

Time and again, we have raised the question of how in the future we ought to treat the artificially intelligent entities (AIEs), such as robots, mindclones, androids, bemans, or any other entity that possesses intelligence, autonomy, and other features that would make them similar to a human being. Furthermore, with human enhancement, and with the prospect of technologies like those aimed at building a robot equipped with a biological brain grown in an incubator or uploading a human brain onto a computer, it is no longer ontologically clear what it is to be human or how we should draw the line between human and nonhuman. This technological scenario confronts us with the ethical problem of inclusion and exclusion: Are the new entities worthy of consideration as moral beings? And, if so, on what basis?

Depending on the way we answer these questions, we will come out with different ways of treating these new entities, thus fundamentally shaping the social environment in which we are going to live in the future.

In this paper I offer an approach to these questions on how to deal with AIEs: My proposal is that we can draw on the insights offered by environmental ethics, and I argue on that basis that AIEs ought to be considered, not as entities extraneous to our social environment, but as forming an integral part of that environment. I proceed in particular from one quite radical strand of environmental ethics, namely, Deep Ecology, which sees all entities as existing in an inter-relational environment: It thus rejects any "firm ontological divide in the field of existence" (Fox) and on that basis introduces principles of biospherical egalitarianism, diversity, and symbiosis. Environmental ethics makes the case that we ought to "include within the realms of recognition and respect the previously marginalized and oppressed" (Gottlieb). I consider whether AIEs can be described along these lines, as somehow "marginalized" or "oppressed," and whether there are grounds for extending to them the kind of recognition that such a description would seem to call for. ↑**Program**↑

Steve Mckinlay
Wellington Institute of Technology
"Evidence, Explanation and Predictive Data Modelling"
Tuesday June 14th, 11:30, Session A

This paper examines the nature and validity of evidence or knowledge generated by predictive data modelling algorithms. In particular I consider Salmon's Statistical Relevance model of scientific explanation and compare this to the way the output of predictive data modelling systems is interpreted. The paper argues that such output cannot be considered evidence in the same sense that most orthodox epistemological approaches might conclude. I conclude by arguing that the only justification that one might have for considering such output as evidence must ultimately rely upon some form of pragmatic justification and that any interventions decisions are ultimately normative. This raises important ethical issues and these are introduced as starting points for further work. ↑**Program**↑

Dimitri Coelho Mollo
King's College London / Humboldt-Universität zu Berlin
"Deflating Representational Content"
Tuesday June 14th, 10:30, Session A

Representation plays a central explanatory role in the cognitive sciences. In order to fulfil this role, theories of representation must meet some requirements: explain how representations come to have the contents they do; give an account of what makes them into representations; and make space for the possibility of misrepresentation. Preferentially, all these requirements should be met by having recourse exclusively to naturalistically acceptable entities and relations, so as to give the notion of representation a respectable place in the scientific worldview.

The notion of representation has traditionally come hand-in-hand with that of computation. The idea that the cognitive system is to be viewed as computational was one of the founding pillars of cognitive science in the 1950s, and is still today a fruitful scientific paradigm. Explaining what computing systems are and in what way the cognitive system is computational, or usefully explained as so being, is another foundational issue in the cognitive sciences.

In this paper, my aim is to provide the outline of a theory of representation able to play the required explanatory role while steering clear from the metaphysical difficulties that plague existing theories. My approach will be deflationary. I will rely on the robust notion of concrete computation provided by the mechanistic view to individuate computational structure as one of the factors that carries the most load in explaining complex appropriate behaviour. Ascription of determinate representational content comes on top of that, and heavily depends on the task at hand and on the particular situation the organism finds itself in.

I will be accepting the invitation made by Piccinini to conjoin existing theories of content, which have traditionally relied on a semantic view of computation, with the non-semantic view of computation provided by the mechanistic account. I take that Structural Representation is a particularly promising candidate for such a treatment. It is a notion of representation which successfully answers Ramsey's 'job description challenge' and which is often at work in empirical research.

The resulting picture is compatible with two metaphysical views on the nature of representation, pragmatism and mild realism. I remain largely neutral on what metaphysical picture to endorse. ↑**Program**↑

Ioan Muntean & Don Howard

University of Notre Dame
"Generalized Moral Agency: Artificial Moral Functionalism and Moral Behaviorism"
Thursday June 16th, 10:30, Session B

One aim of this paper is to generalize the framework of normative agency, beyond the individual human agent, to include the artificial moral agents. Our argument can be read as an attempt to bridge "the abyss" that separates the normative from the descriptive in terms of functional agency rather than in terms of moral/natural properties of actions. We address "universalizability" and "replicability" of moral agency by exploiting some connections between philosophy of mind and ethics. The main frame used here is a version of "moral functionalism," called "dispositional moral functionalism." When virtues are understood as dispositional traits to behavior, the dispositional functionalism constitutes a good candidate for a framework debate on the universalizability and replicability of normative agency. Another consequence of this paper is that once the functionalism and virtue ethics framework are settled, artificial agents are easier to include in normative agency. ↑**Program**↑

Erica Neely

Ohio Northern University
"The Ethics of Choice in Single-Player Video Games"
Tuesday June 14th, 12:30, Session B

Many people have discussed ethics and video games. Thus far the focus has been on two elements: the actions players take within a game and whether certain kinds of games are wrong to create, such as extremely violent video games. Taking Miguel Sicart's work as a key starting place, I emphasize the moral importance of choice with respect to video games. This has several components. Focusing on single-player games, I consider gameplay choices that players make, such as how to complete a particular mission in a game and how the game evaluates and adapts to player choices in a way that attempts to impose and/or encourage a particular moral stance. Much literature focuses specifically on gameplay as the interactive component that makes video games a unique medium. However, I stress that we must also consider the choices designers make when they program a game; these choices can frequently go unnoticed and unquestioned by players. As such, while I agree with those who emphasize that players are capable of moral reflection, I question how often they actually engage in moral reflection. This suggests that there is room for pragmatic concern about how actual players – not ideal players – are affected by video games. ↑**Program**↑

Hyungrae Noh

University of Iowa
"Shannon, Dretske and Millikan on Information and Communication"
Thursday June 16th, 9:30, Session A

Shannon attempts to explain communication without considering informational content. Dretske criticizes Shannon by postulating semantics as the core element of communication in that correlation between the sign and the signed plays a pivotal role in information processes. I show that Millikan's teleosemantics successfully explains how communication can take place without involving the correlation that Dretske mentions. I also show that the blending of Shannon and Millikan's theories successful only to the extent of explaining the information process of communication but, tentatively, not naturalizing intentionality. ↑**Program**↑

Teresa Numerico
Università di Roma
"Big Data"
Thursday June 16th, 12:00, Session B

In this paper I will discuss the future of Big Data as a suite of new methods for social and political research. I will start by tracing a genealogy of the idea that machine can perform better than human beings in managing extremely huge quantity of data, and that the quantity of information could change the quality of the interrogation to those data. It all started in the human-machine interaction as well as in the Artificial intelligence fields.

In the second part of the paper I will analyse Big Data as a social and rhetorical construction of a research policy, claiming in favour of a more detailed account of the consequences for its progressive institutionalization. Without a serious methodological assessment of the changes that these new methods produce in the scientific epistemology of social and political sciences, we risk of underestimating the distortive or uncontrollable effects of the massive use of computer techniques in text analysis, producing situations in which it is very difficult to reproduce the accomplished experiment, and it is arduous to explain the theories that can justify the output of researches. As an exemplification of the problem I will cite the work on emotional contagion published on PNAS in 2014 by data science team leader Kramer and two social scientists from Cornell University.

Until now it is difficult to explore all the consequences of Big Data cognitive project for the idea of human intelligence and for the future of social research methods. The vision that there is no way to manage social data than to follow the results of a machine learning algorithm that works on inaccessible and uncontrollable machines is rather risky and deserve some extra consideration. ↑**Program**↑

Francois Oberholzer & Stefan Gruner
University of Pretoria
"The Notion of 'Information': Enlightening or Forming?"
Thursday June 16th, 10:30, Session A

'Information' is a fundamental notion in the field of artificial intelligence including various sub-disciplines such as cybernetics, artificial life, robotics, etc. Practically the notion is often taken for granted and used naively in an unclarified and philosophically unreflected manner, whilst philosophical attempts at clarifying 'information' have not yet found much consensus within the science philosophical community. One particularly notorious example of this lack of consensus is the recent Fetzer-Floridi dispute about what is 'information' – a dispute which has remained basically unsettled until today in spite of a sequence of follow-up publications on this topic. In this paper our philosophical analysis reveals with reference to Gottlob Frege's classical semiotics that the above mentioned Fetzer-Floridi dispute cannot come to any solution at all, because the two competing notions of 'information' in that dispute are basically synonyms of what Frege had called 'sense' (Sinn) versus what Frege had called 'meaning' (Bedeutung). As Frege had convincingly distinguished sense and meaning very clearly from each other, it is obvious that 'information' understood like 'sense' and 'information' understood like 'meaning' are incompatible and cannot be reconciled with each other. Moreover we also hint in this paper at the often-forgotten pragmatic aspects of 'information' which is to say that 'information' can always only be 'information for somebody' with regard to a specific aim or goal or purpose. 'Information', such-understood, is thus a teleological notion with a context-sensitive embedding into what the late Wittgenstein had called a language-game (Sprachspiel). Shannon's quantified notion of 'information', by contrast, which measures an

amount of unexpected surprise and which is closely related to the number of definite yes-no-questions which must be asked in order to obtain the desired solution of a given quiz puzzle, is not the topic of this paper – although also in Shannon's understanding of 'information' the quiz puzzle scenario, within which those yes-no-questions are asked and counted, is obviously purpose-driven and Sprachspiel-dependent. We conclude our information-philosophical analysis with some remarks about which notion of 'information' seems particularly amenable and suitable for an autonomic mobile robotics project which one of the two co-authors is planning for future work. To separate this suitable notion of 'information' from other ones a new word, 'enlightation', is coined and introduced. ↑**Program**↑

Ugo Pagallo
University of Turin
"The Challenges of Digital Democracy, and How to Tackle Them in the Information Era"
Tuesday June 14th, 9:30, Session B

The aim of the paper is to examine the ways in which scholars commonly address cases of legal disagreement that may potentially concern the broken promises of democracy vis-à-vis the new challenges of the information era. General disagreement may regard the meaning of the terms framing the question, the ways such terms are related to each other in legal reasoning, or the role of the principles that are at stake in the case. By stressing the limits that every theory of justice, or of tolerance, encounters before these cases, the intent is to determine whether the compromises that at times have to be found in the legal and political domains can be deemed as reasonable. ↑**Program**↑

Daniele Porello
Laboratory for Applied Ontology, Institute for Cognitive Sciences and Technologies, ISTC, CNR
"Judgment Aggregation in Relevant Logic"
Thursday June 16th, 12:30, Session A

Judgment aggregation studies the procedures for aggregating a number of sets of logically related propositions provided by heterogeneous agents into a collective set of propositions. The procedures are designed so to represent, as close as possible, the individual views. Judgment Aggregation provides an important framework to discuss the epistemic status of groups of agents since it provides the tools to analyze the dependency of collective information on the individual sources. In this paper, we contribute to the judgment aggregation theory by approaching the aggregation of propositions expressed in relevant logics. ↑**Program**↑

Thomas M. Powers & Jean-Gabriel Ganascia
University of Delaware, Université Pierre et Marie Curie
"Towards a Prescriptive Deontological Non-monotonic Machine Ethics"
Thursday June 16th, 12:30, Session B

Our aims in this paper are twofold: 1) to show the pitfalls of taking the Trolley Problem for machine ethics to be foundational, in particular, by focusing on the implications of its consequentialism; and 2) to give philosophical and practical arguments for a machine ethics that is prescriptive and deontological, and incorporates non-monotonic logic. We demonstrate one approach—inspired by the ethics of Kant—that has these three properties by taking advantage of advances in logic-based Artificial Intelligence, specifically in non-monotonic formalisms, and by programming in a language called Answer Set Programming (ASP). By giving several illustrations of complex ethical situations resolved by deontological ethical arguments programmed

in ASP, we hope to show that it is attractive as an approach, and that it is worthy of further efforts by the machine ethics community. Here, our goal is not to give a complete formal account, with all of the deontological rules illustrated, but rather to provide a framework for future development. ↑**Program**↑

Michele Rapoport
Tel Aviv University
"Dwelling Technologically: Territory and Labor in the Automated Home"
Wednesday June 15th, 10:30, Session B

At the technological forefront as digital devices intended to mediate, improve, extend and even replace the human in the creation of an environment tailored to individuated needs, smart technologies are the next step in technological mechanisms designed to advance the quality of human life. One of the most interesting yet least examined augmented environments that smart technologies facilitate is the home, and indeed smart homes and their technologies are proliferating rapidly. Yet what is the nature of home, and how are homes 'made'? In addressing smart technologies, if and how do these affect, if indeed they do, the nature of dwelling and the making of home? It is on these questions that the following examination will focus as it traces the spatial and performative aspects of home-making, and the constitution of home through human action, behavior and performance, and through the changing nature of labor within the automated domestic realm. ↑**Program**↑

Luca Rivelli
Université Paris 1-IHPST, Università di Padova-FISPPA
"Antimodularity: Pragmatic Consequences of Computational Complexity on Scientific Explanation"
Friday June 17th, 9:30, Session A

This work is mainly concerned with the notion of hierarchical modularity and its use in explaining structure and dynamical behavior of complex systems by means of hierarchical modular models, and proposes a new notion, antimodularity, tied to the possibility of the algorithmic detection of hierarchical modularity. The paper first highlights the pragmatic consequences of hierarchical modularity on the possibility of certain types of scientific explanation of complex systems, that is, systems which, according to a chosen basic description, can be considered as composed of elementary, discrete, interrelated parts. It is stressed that hierarchical modularity is also required by the experimentation aimed to discover the structure of such systems. Networks are chosen as a representative class of complex systems. Algorithmic detection of hierarchical modularity in networks turns out to be a task plagued by the demonstrated computational intractability of the search for the best hierarchical modular description, and by the high computational expensiveness of even approximated detection methods. The proposed notion of antimodularity consists in the unavailability of a modular description fitting the explanatory purposes of the observer, unavailability due either to actual absence of modularity in the system's chosen basic description, or to the impossibility, due to the excessive size of the system under assessment in relation to the computational cost of algorithmic methods, to algorithmically produce a valid hierarchical description. It is stressed that modularity and antimodularity depend on the pragmatic choice of a given basic description of the system, a choice made by the observer based on her explanatory goals. It turns out that antimodularity hinders the possibility of applying at least two well-known types of explanation, mechanistic and functional explanation, which rely on multiple interrelated levels of descriptions and require intelligibility: antimodularity, by impeding the obtainment of a full hierarchical description, negates the possibility of multi-level explanation, thus damaging understandability

of the explanation by allowing only for a description of the system at the lowest level, that of elementary parts, a description which, in sufficiently wide systems, can result unintelligible because of the sheer number of interrelated parts which it describes. Another newly proposed more general notion, explanatory emergence, subsumes all the results about the unfeasibility of certain multilevel explanations and the consequent fading of understandability due to the occurrence of antimodularity, as well as any other case in which a system, for computational reasons, resists understandable explanation because of the excessive computational cost involved in algorithmic methods necessary to obtain the descriptions which are to be employed in explanations. The possible consequences, and the likelihood, of incurring in antimodularity or explanatory emergence in biology and other sciences are finally evaluated, concluding that this eventuality is quite likely, at least in systems biology and related disciplines. ↑**Program**↑

Paul Schweizer
University of Edinburgh
"Computation in Physical Systems: a Normative Mapping Account"
Wednesday June 15th, 11:30, Session A

The relationship between abstract formal procedures and the activities of actual physical systems has proved to be surprisingly subtle and controversial, and there are a number of competing accounts of when a physical system can be properly said to implement a mathematical formalism and hence perform a computation. I develop an account wherein computational descriptions are essentially normative, and where the projected normative criteria vary according to our pragmatic goals. Hence there is no objective or uniform fact to the matter, and I defend the 'anti-realist' conclusion that computational descriptions of physical systems are purely interest relative projections. ↑**Program**↑

Johnny Hartz Søraker
University of Twente
"A Right to be Remembered? Meaningfulness and Permanence in Virtual Environment"
Tuesday June 14th, 12:00, Session B

Although Virtual Environments, including Video Games, may make us experience subjective well-being (commonly referred to as 'happiness') and provide many other benefits, it still seems a stretch to claim that they confer any deep sense of meaningfulness. A life spent doing nothing but playing video games could arguably be said to be a happy life, but is it possible to also refer to it as a meaningful one? After first explaining how I understand the notions of subjective well-being and meaningfulness in this paper, as informed by both empirical research and philosophical accounts, I will go on to argue two related points. First, virtual environments can already be used to bring about extravirtual consequences that provide a sense of meaningfulness – in particular by allowing new forms of altruism. Second, I will also argue that virtual environments can provide a sense of meaningfulness intravirtually, i.e without reference to consequences in the actual world. The latter would, however, require a degree of permanence that is rarely found online. To achieve such permanence, we will either need to recognize a right to be remembered and a corresponding duty on the side of virtual environment providers – or to provide a standard that allows us to safeguard the permanence of our virtual achievements ourselves. I will discuss the feasibility, benefits and problems with both solutions and conclude with normative recommendations for both users and providers. ↑**Program**↑

Ingvar Tjostheim

Norwegian Computing Center
"Telepresence and the Role of the Senses"
Friday June 17th, 10:00, Session A

What is the theoretical basis for the telepresence-phenomenon? Players of videogames can experience telepresence. It is as a subjective experience of being in one place or environment, even when one is physically situated in another place. In this paper we discuss theories of perception, the role of the senses, and the Spinozian belief procedure. A study with the affordance concept is use to empirically investigate how players report sense-reactions in a virtual sightseeing in two cities. ↑**Program**↑

Orlin Vakarelov

Duke University
"Information Media Networks: A Representational Framework for Cognition"
Wednesday June 15th, 9:30, Session A

Cognitive science has struggled to integrate three different theoretical approaches: the information-processing approach, the dynamical system theoretic approach, and the network-based approach. Part of the difficulty comes from the different theoretical machinery used by these approaches. They rely on different formal and mathematical tools for representing the operations and behavior of the cognitive system. We will describe the set of such formal and mathematical tools as the representational framework of the approach. We argue that the Information Medium Network (IMN) representational framework can integrates important theoretical insights from the different approaches. We first argue that IMNs have the necessary elements to serve as a representational framework for cognition. Then we highlight some of the advantages of the representational framework. The advantages are related to two methodological principles: (1) the framework offers a wide possibility space of cognitive architectures; and (2) the framework sets theoretical defaults that favor more inclusive approaches to cognitive architecture. ↑**Program**↑

Sandra Wachter

The Alan Turing Institute
"Privacy - What is it good for? Absolutely nothing?"
Thursday June 16th, 12:00, Session A

There is growing consensus that the Internet poses a threat to privacy and that it is one of the reasons why privacy is either dying or already dead. Nonetheless, academics from various disciplines believe that privacy needs to be protected and call for effective mechanisms to guarantee the free exercise of this human right.

But why engage in a battle that is doomed to be lost?

What is the value of privacy? Do we need privacy for the proper functioning of a democratic state? What is the relationship between privacy and other human rights, such as the right to property, freedom of press and freedom of speech?

The European Court of Human Rights (ECHR) uses democracy as a benchmark when rendering its verdicts. In its jurisprudence the court rules that democracy and human rights are deeply connected. The judgements illustrate the concept of democracy and explain the obligation of states and citizens to protect and promote democratic principles.

In this talk I will describe the concept of democracy according to the ECHR and explain the role of privacy and other human rights in this context. I will argue that the free enjoyment

of fundamental human rights, such as the freedom of religion, the right to choose a profession or the right to vote, are inextricably connected to Internet technologies.

By using this approach, I will first illustrate the value and scope of privacy from a legal perspective and will then move on to conclude that privacy is a necessary precondition for both the proper functioning of a (digital) democratic state and the free enjoyment of other fundamental human rights online. ↑**Program**↑

Anna Wilks
Acadia University
"Robotic Responsibility"
Thursday June 16th, 10:00, Session B

Among the most critical issues for current ethical inquiry are those pertaining to the development of increasingly intelligent and autonomous machines. In this paper I address the question of whether humanoid robots may legitimately be viewed as moral agents capable of participating in the moral community. I defend the view that they can, but in a very qualified sense. Specifically, I contend that they can be viewed as jointly responsible for moral action upon an expanded view of responsibility. The argument of the paper appeals to Immanuel Kant's notions of personhood and moral duty in connection with Margaret Gilbert's idea of joint commitment. I maintain that the combination of these central moral precepts yields a compelling account of how moral responsibility may be attributed to intelligent machines. This account is of considerable utility in a world of increasingly greater interaction between human beings and intelligent machines.

The view I defend is that there is no basis for viewing a robotic being as solely responsible for its behaviour, so long as its behaviour is the result of some joint endeavor or goal – as is usually the case. Because humanoid robots are designed, engineered and operated by persons with an intended end, they may only be viewed as sharing in moral responsibility, since they can share in the intentions, values, actions and commitments of a collective of which they are members, but they cannot be viewed as capable of individual, personal commitments. That is to say, such robotic beings may be viewed as capable of what Gilbert refers to as joint commitment, but not individual commitment. Expanding the concept of moral responsibility in this way establishes a solid basis for machine ethics. ↑**Program**↑

James Williams
University of Oxford
"Three Types of Distraction"
Friday June 17th, 10:00, Session B

When digital technologies distract us, we often regard it as mere annoyance. However, I argue that such a view vastly underestimates the ethical importance of distraction, especially for questions of user autonomy and self-determination. I first clarify the concept of 'distraction' and outline the nature and scope of its effects. I then identify three main types of distraction, which I call Functional, Existential, and Epistemic Distraction. Each type of distraction directs us from a different category of action: from doing, being, and knowing, respectively. For each type of distraction, I discuss examples and explore their ethical implications. Finally, I discuss the ethical obligations that this broader, tripartite view of distraction may present for those who design and deploy information and communication technologies. ↑**Program**↑