

CIVIC AI (CAI): ARTIFICIAL INTELLIGENCE AT THE SERVICE OF SOCIETY

I-LINK Project — Workshop 2

Artificial Intelligence: Human Rights and the Rule of Law

27 June 2026

9:00 a.m. – 1:00 p.m. (CEST)

Venue: Australian Catholic University (ACU)

Organisers

Patrick Keyzer — Professor of Law and Public Policy, Thomas More Law School, Australian Catholic University (ACU), Melbourne, Australia ✉ Patrick.Keyzer@acu.edu.au

Pompeu Casanovas — Research Professor in AI, Ethics and Law at IIIA-CSIC, Barcelona, Spain; Adjunct Professor at La Trobe Law School, Melbourne, Australia; Linked Researcher at IDT-UAB, Barcelona, Spain. pompeu.casanovas@iia.csic.es

Project Manager

Varda Sokolowicz — CSIC Administration. vsokolowicz@iia.csic.es · (+34) 93 580 9570 ext. 431815

Chairs

Patrick Keyzer — Professor of Law and Public Policy, Thomas More Law School, Australian Catholic University (ACU), Melbourne, Australia ✉ Patrick.Keyzer@acu.edu.au

An additional Chair

Links:

CAI Workshop Zoom Meeting

[Tiva, please insert Zoom link here]

Overall Objective

This Second Workshop explores the relationship between AI, Law, Ethics, and Civic Technologies (Civic Tech), and is specifically concerned with the interaction between AI, Human Rights and the Rule of Law. These CSIC projects are aimed at creating research networks. In accordance with the proposal, the programme pursues three interrelated aims: (i) the design and governance of social ecosystems through advanced artificial intelligence techniques — including electronic institutions, multi-agent systems, and Large Language Models (LLMs); (ii) the alignment of such systems with the ethical and legal imperatives of the rule of law, encompassing the protection of human rights against social, ethnic, and gender discrimination, the regulation of emergent financial markets, and the cultivation of resilient and sustainable human–machine interactions; and (iii) the consolidation of a scientific network committed to the development of intelligent regulatory frameworks.

The planned Workshops examine the operational convergence between Civic Artificial Intelligence (CAI) and Civic Technologies — digital instruments conceived to foster open, transparent, and accountable governance. Grounded in the conviction that artificial intelligence must remain at the service of society, they bring together research on human-centred electronic institutions, LawTech, RegTech, and computational law, with a view to addressing the normative, ethical, and policy dimensions of increasingly complex digital environments.

Specific Objective

This second Workshop is specifically concerned with the interaction between AI, Human Rights and the Rule of Law.

The format is deliberately flexible, accommodating both focused contributions — including work in progress by doctoral researchers — and broader presentations of consolidated research projects. This first gathering is intended, above all, as an opportunity for mutual orientation: each team is invited to articulate the distinctive contours of its research agenda, enabling all partners to develop a clear and comparative understanding of one another's perspectives and priorities. The following questions may serve as a structural guide: (i) What are the fundamental research questions and objectives pursued, and what results have already been obtained? (ii) Upon what theoretical and methodological foundations does the work rest? (iii) What analytical and empirical methods are employed? (iv) What are the projected next steps?

These questions are offered as suggestions only. Each participant is entirely free to determine the form and emphasis they consider most appropriate for presenting their research — including, where relevant, the sharing of results of publishable quality. In that event, a double-blind peer-review process will take place after the Workshop.

List of Topics

The workshop topics, grouped in four thematic clusters, include but are not limited to:

1. AI Core for Computational Law and Regulatory Systems

- Automated Compliance: architectures for monitoring adherence to the rule of law in private and public sectors.
- RegTech and LawTech innovations: AI-driven tools for legal reasoning, contract analysis, and regulatory reporting.
- Smart Electronic Institutions: design and formalisation of normative multi-agent systems.
- Algorithmic Governance: mechanisms for enforcing transparency and accountability in automated decision-making.
- Values and principles models: modelling.

2. AI-Driven Civic Technologies and Public Administration

- AI for Open Government: enhancing Open Government Partnership (OGP) objectives through intelligent data analysis.
- Participatory Platforms: integrating LLMs and agents into participatory budgeting and public debate forums.
- Service Delivery Optimisation: using AI to improve responsiveness and equity in public service distribution.
- Digital Democracy: tools for mobilising citizen engagement and combating corruption in decentralised environments.

3. AI Interplay with Ethics, Human Rights, and Social Ecosystems

- Algorithmic Fairness: detection and mitigation of social, ethnic, and gender discrimination in civic AI systems.
- Human-Centric Interfaces: designing resilient human–machine interaction (HMI) for public services.
- Data Sovereignty and Privacy: reconciling open-data initiatives with individual privacy rights.
- Resilience and Sustainability: long-term viability of AI-mediated social ecosystems in changing political landscapes.
- Online Dispute Resolution: technology and AI for automated and semi-automated conflict resolution.

4. Technical Foundations of Civic AI

- LLMs in Governance: fine-tuning and prompt engineering for legal and civic applications.
- Multi-Agent Simulations: modelling policy impacts and social dynamics before implementation.
- Trust and Reputation Models: computational approaches to justifying official actions and assigning responsibility.
- Agentic AI Governance: process models within legal constraints.

Programme

All times are CEST. Sydney / Melbourne / Brisbane are 8 hours ahead (05:00 p.m. – 09:00 p.m.). Each presentation is allocated 20 minutes, inclusive of questions. The programme is organised in four thematic clusters, two per day.

Friday, 26 June

Clusters I and II — Computational Law and Regulatory Systems · Civic Technologies and Public Administration

Opening Session	
09:00 – 09:05	Welcome and presentation of the agenda — Professor Patrick Keyzer (ACU)
09:05 – 09:15	Brief summary of the CAI i-Link Project — Pompeu Casanovas (IIIA-CSIC / La Trobe / UAB)
09:15 – 09:35	Ozana Olariu, University of Granada, John Zeleznikow (LTU). <i>Human Rights-Compliant Artificial Intelligence - Regulatory Frameworks for Upholding Democracy and the Rule of Law in the Digital Age</i>
09:35 – 09:55	Discussion.
09:55 – 10:15	Thu Nguyen, Simeon J. Simoff, David Arness, Quang Vinh Nguyen, <i>The Mindful Political Economy of AI: Emancipatory Approaches</i>
10:15 – 10:35	Discussion
10:35 – 10:55	Patrick Keyzer (ACU), Mustafa Hashmi (ACU), Pompeu Casanovas (IIIA-CSIC). <i>Theorising Access to Justice for People in Closed Environments</i>
10:55 – 11:15	Discussion
11:15 – 11:35	
11:35 – 11:40	Short break
11:40 – 12:00	
12:00 – 12:20	
12:20 – 12:40	

12:40 – 13:00	
---------------	--

Abstracts

Abstracts are presented in the order in which the contributions appear in the programme.

Patrick Keyzer (ACU), Mustafa Hashmi (ACU), Pompeu Casanovas (IIIA-CSIC). *Theorising Access to Justice for People in Closed Environments*

For human rights to work, they need to be accessible, or open. In this “White Paper” we position legal ecosystems as a way to implement the Mandela Rules (the Standard Minimum Rules for the Treatment of Prisoners). Adopting a legal-ethical value chain model, we suggest how compliance can be modelled, enabling implementation. Two caveats frame this project. First, its realisation does not depend on technological breakthroughs yet to come: the computational tools required — large language models, formal-logic engines, knowledge-graph architectures, and explainable-AI frameworks — are already mature and widely deployed, so the task is one of adaptation and integration rather than invention. Second, artificial intelligence and automated data analysis are to be understood as the final layer of the ecosystem rather than its starting point: they can legitimately be built only once a solid framework for controlling, validating and monitoring their outputs is already in place. We argue that the model we propose here would be particularly beneficial in closed environments; for prisoners, prison administrators and prisoner advocates. Our aim is to achieve “open rights” and contribute to what can be termed Civic AI, a specific branch of Civic Technologies focused on the use of AI tools to enhance and implement regulatory models and ethical and legal ecosystems. The deployment of artificial intelligence to state the law and resolve disputes in real time creates extraordinary opportunities. It also creates risks which are considered in this paper. In this sense, the smart legal ecosystem we describe is a governance project that happens to rely on technologies already at hand, but is oriented and capable of creating new formal languages to extract rules from norms and establish the necessary framework for the generation of the legal ecosystem that is needed for the implementation of the Mandela rules.

This is a White paper. It does not report an experiment, a proof of concept, or a deployed use case. Its contribution is conceptual and architectural. We set out the benefits, and equally the challenges, of three tasks: (i) modelling the Mandela Rules so as to build a genuinely useful Civic AI tool; (ii) building a smart legal ecosystem that includes all of its stakeholders — not only prisoners and their advocates, but also prison officers, inspectors, judges, administrators, police and governments; and (iii) modelling across at least three levels of abstraction — micro, meso and macro — together with the embedded contexts within which norms acquire their meaning. Deontic event-calculus is used only at the close of the paper, to formalise a small number of Rules by way of illustration, and to expose the difficulties of formalisation as candidly as its possibilities.

Ozana Olariu, University of Granada, John Zeleznikow (LTU). *Human Rights-Compliant Artificial Intelligence - Regulatory Frameworks for Upholding Democracy and the Rule of Law in the Digital Age*

AI offers unprecedented opportunities alongside profound challenges for human rights, democracy, and the rule of law. We explore the necessity of integrating human rights principles into AI development and governance to mitigate risks such as algorithmic bias, discrimination, and threats to fundamental freedoms.

We highlight how AI's pervasive influence requires a robust, multi-level regulatory framework grounded in international and regional human rights instruments.

The paper examines the European Union's AI Act and the Council of Europe's Framework Convention on AI, Human Rights, Democracy, and the Rule of Law, emphasizing their complementary approaches and highlights their pitfalls, in the broader regional regulatory context.

The analysis underscores the global imperative for human rights centric AI regulation that prioritizes human dignity, equality, and accountability.

Thu Nguyen, Simeon J. Simoff, David Arness, Quang Vinh Nguyen, *The Mindful Political Economy of AI: Emancipatory Approaches*

This paper adopts a conceptual research design, based on an integrative literature review, to examine how power, responsibility, and attention are organised within contemporary AI systems as an emerging "AI empire". Drawing on this analysis, the paper develops a theoretically grounded framework to deepen understanding of the unfolding contexts of AI and the prospects for better practices. The paper makes two key contributions. First, it proposes the Mindful Political Economy of AI (MPEA) as a four-level framework that links individual awareness, organisational practices, institutional governance, and political-economic structures. This work reveals limitations in many responsible AI approaches, which often prioritise technical and ethical issues while paying insufficient attention to underlying power relations. It highlights attention and awareness as central processes through which AI shapes human subjects in the AI Empire. Second, the paper explains how AI systems shape human attention, awareness, and agency. It shows how AI influences how people think, make decisions, and engage with the AI Empire-mediated environments. It provides new insights into how AI structures human experience and suggests avenues for broader social outcomes in confronting the Empire of AI.

Programme Committee

- Pompeu Casanovas (IIIA-CSIC, Barcelona, Spain)
- Louis de Koker (La Trobe University, Melbourne, Australia)
- Marina Echebarría Sáenz (UVa, Valladolid, Spain)
- Enrico Francesconi (CNR, Florence, Italy)
- Carles Górriz López (UAB, Barcelona, Spain)
- Guido Governatori (CQU, Brisbane, Australia)
- Mustafa Hashmi (ACU, Melbourne, Australia)
- Patrick Keyzer (ACU, Melbourne, Australia)
- Ho-Pun Lam (Xi'an Jiaotong-Liverpool University, XJTLU, China)
- Lissette Lemus del Cueto (IIIA-CSIC, Barcelona, Spain)
- Ramon López de Mántaras (IIIA-CSIC, Barcelona, Spain)
- Tarik Nesh-Nash (OGP, Tangier, Morocco)
- Pablo Noriega Blanco-Vigil (IIIA-CSIC, Barcelona, Spain)
- Nardine Osman (IIIA-CSIC, Barcelona, Spain)

- Carmen Pastor Sempere (UA, Alicante, Spain)
- Enric Plaza Cervera (IIIA-CSIC, Barcelona, Spain)
- Marta Poblet (Data Tank, Brussels, Belgium)
- Ante Prodan (University of Sydney, Australia)
- Víctor Rodríguez-Doncel (UPM, Madrid, Spain)
- Antonino Rotolo (University of Bologna, Italy)
- Giovanni Sartor (University of Bologna / EUI, Bologna / Florence, Italy)
- Carles Sierra García (IIIA-CSIC, Barcelona, Spain)
- Simeon J. Simoff (WSU, Sydney, Australia)
- Varda Sokolowicz (IIIA-CSIC, Barcelona, Spain)
- Marina Stefanova (Sofia University, Bulgaria)
- Esther Zapater Duque (UAB, Barcelona, Spain)
- John Zeleznikow (La Trobe University, Melbourne, Australia)

Presenters

- Adam, Rani — Western Sydney University, Australia ✉ 22104050@student.westernsydney.edu.au
- Casanovas, Pompeu — IIIA-CSIC, Barcelona, Spain / La Trobe University, Melbourne, Australia / IDT-UAB, Barcelona, Spain ✉ pompeu.casanovas@iiia.csic.es
- Catchpoole, Daniel R. — University of Technology Sydney (UTS), Australia (Professor, Biomedical Data Science Lab) / Western Sydney University (Adjunct) ✉ daniel.catchpoole@health.nsw.gov.au
- Coluccino, Delia Maria — LAST-JD 39th cycle, University of Bologna, Italy / UAB, Barcelona, Spain ✉ — deliamaria.coluccin2@unibo.it
- de Koker, Louis — La Trobe Law School, La Trobe University, Melbourne, Australia ✉ l.dekoker@latrobe.edu.au
- Governatori, Guido — Central Queensland University (CQU), Brisbane, Australia ✉ g.governatori@cqu.edu.au
- Hashmi, Mustafa — La Trobe Law School, La Trobe University, Melbourne, Australia / Australian Catholic University (ACU), Melbourne, Australia ✉ pakdost76@gmail.com
- Huang, Haoyu — Xi'an Jiaotong-Liverpool University (XJTLU), China ✉ — to complete —
- Kennedy, Paul J. — School of Computer Science, University of Technology Sydney (UTS), Australia ✉ Paul.Kennedy@uts.edu.au
- Keyzer, Patrick — Thomas More Law School, Australian Catholic University (ACU), Melbourne, Australia ✉ Patrick.Keyzer@acu.edu.au
- Khan, Hassan — Western Sydney University, Australia (Social Robot Interaction Group) ✉ — Hassan.Khan@westernsydney.edu.au
- Lam, Ho-Pun — Department of Computing, Xi'an Jiaotong-Liverpool University (XJTLU), China ✉ Ho.Lam@xjtlu.edu.cn
- Ramon López de Mántaras — IIIA-CSIC mantaras@iiia.csic.es
- Mubin, Omar — School of Computer, Data and Mathematical Sciences, Western Sydney University, Australia ✉ O.Mubin@westernsydney.edu.au

- Nesh-Nash, Tarik — Open Government Partnership (OGP), Tangier, Morocco ✉ tarik.nesh-nash@opengovpartnership.org
- Nguyen, Quang Vinh — School of Computer, Data and Mathematical Sciences, Western Sydney University, Australia ✉ q.nguyen@westernsydney.edu.au
- Noriega, Pablo — IIIA-CSIC, Barcelona, Spain ✉ pablo@iia.csic.es
- Osman, Nardine — IIIA-CSIC, Barcelona, Spain ✉ nardine@iia.csic.es
- Pisano, Giuseppe — Legal Machine Lab, CIRSFD, University of Bologna, Italy ✉ g.pisano@unibo.it
- Plaza, Enric — IIIA-CSIC, Barcelona, Spain ✉ enric@iia.csic.es
- Poblet, Marta — The Data Tank, Brussels, Belgium / RMIT University, Melbourne, Australia (Adjunct Professor) ✉ marta.poblet@datatank.org
- Qu, Zhonglin — School of Computer, Data and Mathematical Sciences, Western Sydney University, Australia ✉ z.qu@westernsydney.edu.au
- Raulino, Thiago — University of Bologna (PhD), Italy ✉ thiago.raulino@unibo.it
- Rodríguez-Doncel, Víctor — Ontology Engineering Group (OEG), Universidad Politécnica de Madrid (UPM), Spain ✉ vrodiguez@fi.upm.es
- Romero, Margarida — IIIA-CSIC, Barcelona, Spain / Université Côte d'Azur, Nice, France, ✉ margarida.romero@iia.csic.es
- Sartor, Giovanni — CIRSFD-Alma AI, University of Bologna, Italy / European University Institute (EUI), Florence, Italy ✉ giovanni.sartor@unibo.it
- Sierra, Carles — IIIA-CSIC, Barcelona, Spain ✉ sierra@iia.csic.es
- Simoff, Simeon J. — School of Computer, Data and Mathematical Sciences, Western Sydney University, Australia ✉ s.simoff@westernsydney.edu.au
- Sudano, Giulia — LAST-JD 41st cycle, University of Bologna, Italy ✉ giulia.sudano2@unibo.it
- Vagnoni, Simone — University of Bologna (PhD), Italy / Universidad Politécnica de Madrid (UPM), Spain ✉ simone.vagnoni3@unibo.it
- Valizadeh, Jaber — Western Sydney University, Australia J.Valizadeh@westernsydney.edu.au
- Vallès, Núria — IIIA-CSIC, Barcelona, Spain ✉ nuria.valles@iia.csic.es
- Zeleznikow, John — La Trobe Law School, La Trobe University, Melbourne, Australia / Joseph Bell Centre, University of Edinburgh Law School, UK ✉ j.zeleznikow@latrobe.edu.au
- Zhang, Dongmo — School of Computer, Data and Mathematical Sciences, Western Sydney University, Australia ✉ d.zhang@westernsydney.edu.au