

SEPTEMBER 27TH, 2026 VIRGINIA TECH, WASHINGTON, D.C.

The First Workshop on **Sovereign AI** for Collaborative and Pluralistic AI Ecosystems

As today's AI ecosystem is shaped by models, data, and compute concentrated in a few global institutions, discussions of "sovereign AI" are often reduced to a question of infrastructure. This workshop asks a deeper question: who holds the power to decide what stories AI tells, whose communities it represents, and whose values it is aligned to?

i. **Overview**

Sovereignty is often equated with owning one's own compute and models, but compute and models are ultimately the means of sovereignty, not its ends. This workshop convenes the AI, HCI, and law & policy

communities around a shared question: how can technical capacity, interaction design, and governance institutions together keep AI systems contestable and answerable to the communities they represent?

ii. Topics of Interest

01 Safety and Alignment for Agentic & Physical AI

How do we ensure that increasingly autonomous and embodied AI systems remain safe, aligned, and under meaningful human control, especially when they are adapted or deployed outside the contexts in which they were built?

› Possible topics 9

02 Open Models, Data Sovereignty, and Responsible Deployment

How can open models and local data governance move beyond the defensive goal of “reducing dependence” to become positive conditions for cultural and narrative agency? How can communities build on global models while retaining control over how their data, corpora, and knowledge are collected, interpreted, and used?

› Possible topics 4

03 Human Agency & Trustworthy Human-AI Interaction

How can people effectively understand, oversee, and intervene in AI systems without losing their own agency?

› Possible topics 10

04 Inclusive AI for Underserved Communities

How can people from different technical, cultural, and professional backgrounds help build AI that better serves diverse languages, communities, and public needs?

› Possible topics 8

05 Democratic Resilience and Multi-Stakeholder AI Governance

How can governments, researchers, industry, and the public meaningfully participate in governing AI and holding it accountable? And when “the local community” is itself plural, even internally contested, who has the legitimacy to speak on its behalf?

› Possible topics 11

iii. **Confirmed Speakers**

Confirmed researchers across human computation, social complexity, computational social science, and AI governance.



Isabel Hou

Deputy Minister, Ministry of Digital Affairs
(Taiwan)

TALK TITLE:

Sovereign AI through Collaboration:
Taiwan's Policy Strategy and AI
Governance in Practice



Pietro Michelucci

President, Human Computation Institute
Visiting Professor, Cornell University,
Meinig School of Biomedical Engineering

TALK TITLE:

Who Pulls the Strings? Purpose, Power,
and Human Agency in AI



Senjuti Dutta

Postdoctoral Fellow, Center for the
Dynamics of Social Complexity (DySoC),
University of Tennessee



Giordano De Marzo

Postdoctoral Researcher, Computational
Social Science Lab, Department of Politics

and Public Administration, University of
Konstanz

TALK TITLE:

Emergent behaviours in AI
agents societies



Nicholas Caputo

Assistant Professor of Frontier AI
Governance, Johns Hopkins School of
Government and Policy

TALK TITLE:

Sovereignty, Alignment, and Pluralism

iv. **Tentative Outline**

Keynote order and exact timing are still being finalized with our speakers.

PROGRAM

9:00–9:05	Opening
9:05–9:40	Keynote Speaker: Deputy Minister Isabel Hou 協作共創的主權 AI：台灣政策佈局與 AI 治理實踐 <i>Sovereign AI through Collaboration: Taiwan's Policy Strategy and AI Governance in Practice</i>
9:40–10:15	Keynote Speaker: Dr. Pietro Michelucci <i>Who Pulls the Strings? Purpose, Power, and Human Agency in AI</i>
10:15–11:15	Oral Presentation Sessions
11:15–11:55	<u>Hacking Sovereign AI</u> Deliberative democracy with Murmi.
11:55–12:25	Keynote Speaker: Prof. Nicholas Caputo <i>Sovereignty, Alignment, and Pluralism</i>

12:25–12:55

Keynote Speaker: Prof. Giordano De Marzo

Emergent behaviours in AI agents societies

12:55–

Closing

Interested in attending?

Registration is handled through the main conference. Register for HCOMP 2026 and select this workshop.

[Register for HCOMP 2026](#) →

v. **Hacking Sovereign AI**

Join the discussion. Shape the report.

What should a Sovereign AI know, and how should it respond?

Participants will define the public-interest questions a Sovereign AI must address, the evidence and local context its answers require, and the value conflicts those answers must keep visible.


Deliberation with Murmi

Participants from AI, HCI, law, public policy, and other fields will examine each focal question through three lenses:

- A Which information and sources does a good answer require?
- B Which local contexts, stakeholders, and value judgements must the model recognise?

C

Which trade-offs or value conflicts should remain visible rather than being simplified into one answer?



Murmi converts discussion into claims participants can vote on, revealing opinion clusters, tensions, and possible common ground.

How it works

1

Explore

Review questions grounded in public issues, news, and competing perspectives.

2

Select

Choose 1-2 questions that are representative, contested, or dependent on local context.

3

Deliberate

Define what a good answer requires: evidence, stakeholders, values, and trade-offs.

4

Map perspectives

Use Murmi to surface claims, consensus, and disagreement, then review and correct the map.

5

Build together

Turn the discussion into answer criteria and inputs for dataset design and evaluation.

Why participate?

Sovereign AI should reflect more than the priorities of model developers. Your expertise helps define what good, locally grounded answers look like.

Evaluation framework

Criteria for knowledge, context, uncertainty, values, and trade-offs.

Workshop report

A record of focal questions, consensus, disagreement, and recommendations.

Research material

Inputs for the QA Dataset, benchmark rubric, papers, and follow-on studies.

If the discussion yields meaningful insights, we plan to publish a post-workshop report on arXiv, with participants optionally acknowledged as contributors.



METHODOLOGICAL FOUNDATION

Learning from vTaiwan

This session adapts vTaiwan's approach to open public-policy deliberation for the design and evaluation of Sovereign AI. vTaiwan provides the participatory model, while Murmi helps the room surface claims, map agreement and disagreement, and refine the results together.

vTaiwan brings government, experts, businesses, and citizens into a shared process that combines online participation with in-person deliberation. Rather than making policy decisions, it helps diverse participants clarify competing positions and develop recommendations that can inform policy.

PROCESS

Frame issues → Gather input → Deliberate → Develop recommendations

Openness Transparent information and an accountable process

Collaboration Broad participation and constructive dialogue

Co-creation Shared problem-solving and practical outcomes



A CASE FROM TAIWAN

Taiwan Public Media Living Memory

As generative AI increasingly decides how the world is told, media organizations are among the first to sense what is at stake: whose language, whose history, and whose memory gets written into the next generation's knowledge base will shape how a society is understood — and who gets to decide.

[Read more](#) ›

vi. Call for Papers

PAPER SUBMISSION DEADLINE

August 31

2026

[Submit via EasyChair](#) →

ACCEPTANCE NOTIFICATION

September 6, 2026

WORKSHOP DAY

September 27, 2026

DEADLINES CLOSE

11:59pm Anywhere on Earth

This workshop invites researchers and practitioners to rethink sovereign AI through the lens of human-AI collaboration, interaction, and governance: when global models inevitably shape local lives, how can communities retain the capacity to define their own narratives, values, and collective futures?

We welcome contributions that advance, critique, or reimagine sovereign AI through three complementary paths including **technical** (models, data, safety, and evaluation), **design** (interaction, oversight, and intervention), and **institutional** (law, policy, and governance). We especially encourage work that crosses these boundaries. We particularly welcome submissions from researchers and practitioners in **AI, HCI, law and public policy, government, journalism, civil society, open-source communities**, and related fields.

SUBMISSION TYPES

Extended abstracts, work-in-progress reports, or real-world case studies.

LENGTH AND FORMAT

2-4 pages (excluding references) using the ACM conference format.

PRESENTATION

Accepted submissions will be presented in **oral or poster sessions** and will serve as the foundation for the workshop's interactive discussions and collaborative activities.

ACM TEMPLATES

See [General Submission Instructions](#) for the official ACM templates and formatting policies.

Papers in our workshop are non-archival; accepted work may be submitted or published elsewhere without restriction.

vii. Organizers

Workshop Organizers



Dr. Lun-Wei Ku



Wan-Jhen (Crystal) Wu

Research Fellow, Institute of Information Science, Academia Sinica,
Taiwan

lwku@iis.sinica.edu.tw



Joshua C. Yang

Postdoctoral Researcher, ETH Zurich · Incoming Postdoctoral Fellow,
MIT Laboratory for Information & Decision Systems

joshiyang@mit.edu

Technical Program Manager, Taiwan Public Media NewsLLM Project,
Academia Sinica, Taiwan

apokoios@as.edu.tw



Organizing Committee



Teresa Copple

University of California, Los Angeles



Yu-Hsuan Hsiao

National Tsing Hua University



Shi-Wei Dai

National Taiwan University



Poren Chiang

LL.M., UCLA School of Law

Sovereign AI Workshop · Co-located with [HCOMP 2026](#)