2026 Winter School on AI and Justice Report

Executive Summary

1) Introduction

Last week, Sciences Po Law School and Sorbonne Université jointly hosted the **Winter School on AI and Justice**, an intensive, interdisciplinary programme examining how AI systems increasingly **mediate classification, prediction, and decision-making** across society—and what this shift means for **justice**, understood broadly (beyond courts) as the conditions under which people access rights, opportunities, protections, and remedies.

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

Building on the School's framing that **"justice cannot be reduced to a metric and law cannot be reduced to a checklist,"** the programme pursued two core objectives:

- **Analytical:** equip participants to critically understand how algorithmic systems reshape legal phenomena, governance, and institutional practices.
- **Practical ("critique to craft"):** develop participants' capacity to translate principles into workable legal, technical, and institutional solutions—culminating in a hands-on hackathon.

Participants. The cohort comprised **25 students**:

- **16 from Sciences Po** (14 Law School; 1 School of Public Affairs; 1 Paris School of International Affairs)
- **9 from Sorbonne Université**, primarily students in **computer science**, with additional backgrounds in **mathematics** and **sociology**. This composition supported the Winter School's aim of sustained dialogue between legal reasoning, technical knowledge, and social-science perspectives.

2) Project Information

The Winter School and hackathon were organized in the context of **AI-NODES (Normative and Distributive Effects of AI)**, co-led by:

- **Sciences Po Law School:** Séverine Dusollier, Rebecca Mignot-Mahdavi, Raphaële Xenidis
- **Sorbonne Cluster for AI (SCAI), Sorbonne Université:** Arnaud Latil

AI-NODES is part of **PostGenAI@Paris**, funded by the **Agence Nationale de la Recherche (ANR)**. The event was delivered thanks to the essential work and support of research assistant and postdoctoral researchers **Marta Arisi, Simona Ramos, Klaudia Klonowska, and Nicolas Malevé**, and in partnership with **Clifford Chance**.

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

3) Program Overview

The Winter School unfolded over five days, moving from foundational questions to applied governance and compliance challenges:

- **Day 1 (Sciences Po):** Opening sessions on whether AI can support fair decisions, and the broader **social implications** of AI systems.
- **Day 2 (SCAI, Sorbonne Université):** Deep dive into **fairness operationalization**, tensions between fairness metrics, and emerging compliance approaches (including **verifiable proofs** for black-box systems). The day also addressed **algorithmic discrimination** in legal terms and featured a session on **free and open-source AI** ecosystems.
- **Day 3 (Sciences Po):** Focus on AI in **jurisdictional and quasi-jurisdictional processes**, followed by "HackPrep" sprint exercises to prepare for the hackathon's interdisciplinary work.
- **Day 4 (Clifford Chance):** Full-day **legal hackathon**, co-organised with Clifford Chance and with the precious support of Partners Dessislava Savova, Alexander Kennedy and Frédérick Lacroix, with the presence and insights from the **Défenseur des Droits**.
- **Day 5 (Google AI Lab):** Practitioner-oriented sessions on **explainability** (with insights from Google DeepMind) and legal issues around AI models, including user protection and adjacent IP questions.

Overall, the programme combined:

- conceptual grounding (justice, fairness, institutional power),
- legal frameworks (equality/discrimination, accountability, compliance),
- technical understanding (model behavior, bias mechanisms, explainability), and
- real-world constraints (documentation practices, audits, operational risk, procurement, and deployment).

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

4) Speakers and key contributors

The Winter School brought together an interdisciplinary set of contributors spanning **computer science, law, policy, regulatory practice, and critical socio-technical research**. Across sessions, speakers collectively addressed (i) the **limits of fairness metrics**, (ii) the translation of bias into **legal equality/non-discrimination** claims, (iii) **governance and compliance** pathways (documentation, audits, and institutional responsibilities), and (iv) practical constraints around **explainability** and real-world deployment.

- **Agathe Balayn** (Microsoft Research, New York City) — research on harms of deployed ML systems, ML supply chains, and practical remedies informed by qualitative inquiry.
- **Stefania di Stefano** (CNAM) — regulation of online content moderation; implementation and fundamental-rights implications of the **EU Digital Services Act**.
- **Seda Gürses** (TU Delft; KU Leuven affiliate) — privacy engineering and protective optimization technologies; infrastructures, social justice, and political economy of computing.
- **Klaudia Klonowska** (Sciences Po) — AI decision-support systems and human-machine interaction; international law and emerging technologies in security/defence.
- **Fabian Lütz** (University of Ottawa) — AI regulation, gender equality and non-discrimination; EU equality-law lens on algorithmic discrimination and oversight.
- **Nicolas Malevé** (Sciences Po) — artist and researcher on socio-technical networks of AI, data practices (including annotation), and controversies around generative image platforms.
- **Nicolas Maudet** (Sorbonne Université, LIP6) — collective and distributed decision-making; computational social choice, argumentation and deliberation; explanation in decision-aiding.

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

- **Fabio Pianese** (Nokia Bell Labs) — research leadership on infrastructure systems design and deployment; perspective on complex technical systems and operational realities.
- **Simona Ramos** (SCAI, Sorbonne Université) — formalising legal requirements into verifiable technical architectures; compliance gaps in evolving algorithmic systems.
- **Bruna Sellin Trevelin** (Hugging Face) — legal counsel at the intersection of AI, IP, privacy and competition.
- **Claire Strugala** (French Court of Cassation) — digital dissemination of case law, open data, and institutional transformation of courts through digital tools.
- **Fabien Tarissan** (CNRS; ENS Paris-Saclay; CNIL member) — networks, recommendation systems, diversity and discriminatory effects; public decision-making implications.
- **Nathalie Vandystatd** (European Commission, **European AI Office**) — stakeholder engagement and policy communication; experience coordinating AI Act negotiations.
- **Victor Fabre** (Google France) — litigation and regulatory counsel; perspective on enforcement dynamics and compliance in practice.

Key contributors

- **Legal hackathon with the presence and insights from the Défenseur Des Droits (Bilkiss Omarjie, Elia Chanet, Gabrielle Du Boucher).**
- **Benoit Rottembourg (INRIA):** Contributed to the hackathon by **presenting and framing the case study** and supporting the teams' engagement with the scenario's technical and socio-technical dimensions (fraud detection constraints, bias signals, and operational trade-offs).
- **Dessislava Savova, Alexander Kennedy and Frédérick Lacroix, Clifford Chance.**
- **Victor Fabre, Google.**

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

5) Hackathon Details

Case study (scenario summary)

The hackathon centered on a simulated but realistic dispute involving a high-risk AI system used in banking onboarding and fraud prevention. The provider **Mycheat.ai** (fictitious) supplies a fraud- and risk-scoring model adopted by a large share of European banks. The case introduced a classic governance tension:

- signals that improve fraud detection may be **unevenly distributed across groups**,
- resulting in disparate error rates (notably **higher false positive rates for older applicants**),
- while bias mitigation can reduce disparity but may weaken **KYC/AML effectiveness**.

A rejected applicant ("the Victim") receives only a generic message ("Risk score too high") and suspects discrimination after observing similar refusals among peers. A whistleblower data scientist ("Benoit") alleges internal awareness of fairness issues, saves internal materials (including a partial dataset extract and evaluation documents), and contacts a consumer association and journalist; the regulator is alerted.

Working format

Teams were assigned stakeholder roles (Provider, Bank/Deployer, Victim, and Regulator) and tasked with producing a **legal strategy** consistent with their role and incentives—balancing:

- individual rights and access to explanation,
- allocation of responsibility across provider/deployer,
- documentation and dataset governance,
- trade secrets and confidentiality,
- evidence and burden of proof, and
- systemic risk given broad market adoption.

Evaluation basis

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

The judging relied on the attached **Hackathon Evaluation Criteria**, which rewarded:

1. **Legal analysis and argumentation** (substantive and procedural coherence)
2. **Understanding of technology and algorithmic decision-making** (ability to explain system behavior and bias pathways)
3. **Interdisciplinary integration** (law + tech + ethics + societal impacts)
4. **Responsiveness and critical engagement** (handling counterarguments)
5. **Clarity and persuasiveness of presentation** (structure, evidence, accessibility)
6. **Team collaboration and role realization** (cohesion, balanced participation)

Jury

- Alix Cravero, Coordination Manager – Financial Security Division, Société Générale
- Ghislain Houssel, Head of legal à Qonto
- Gabrielle Du Boucher, Défenseur Des Droits
- Bilkiss Omarjie, Défenseur Des Droits
- Frédérick Lacroix, Partner, Clifford Chance
- Fabien Tarissan, computer science researcher at the CNRS and a professor at the École Normale Supérieure Paris-Saclay
- Beatriz Botero Arcila, Assistant Professor, Sciences Po

Winning team (and prize process)

The winning team was the team of the Regulator:

- **Salim Cochereau** (Sciences Po Law School)
- **Ève Le Riche** (Sciences Po Law School)
- **Cécile Freiin von Roenne** (Sciences Po Law School)
- **Issa Ka** (Sorbonne Université, Computer Sciences)

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

- **Yuxuan Pei** (Sorbonne Université, Sociology and Mathematics)

All hackathon teams were intentionally composed to maximize interdisciplinarity (law + computer science and allied disciplines). The winning team's members are now entering a **selection process**, at the end of which **one participant will be offered an internship at Clifford Chance**.

6) Feedback and Reflections

Participant feedback is currently being collected. While analysis is ongoing, the programme design strongly suggests several likely "signal areas" to capture in the final feedback synthesis:

What worked particularly well (dimensions to confirm via survey data):

- **Interdisciplinary pedagogy:** sustained interaction between legal and technical approaches, especially during hack prep and the hackathon.
- **Applied learning:** the hackathon case study effectively forced participants to confront real governance trade-offs (fairness vs fraud detection; transparency vs trade secrets; provider vs deployer responsibility).
- **Institutional diversity:** exposure to academic, regulatory, legal practice, and industry perspectives (including sessions hosted at partner institutions).

Key learning tensions surfaced (to structure reflections):

- **Metrics vs rights:** fairness metrics do not map neatly onto legal equality and non-discrimination standards; legal remedies require narratives, evidence, and procedural safeguards.
- **Explainability in practice:** "meaningful explanation" demands both technical capacity (what can be explained) and institutional readiness (how it is communicated and contested).

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

- **Accountability allocation:** complex value chains (provider + deployer + auditors + regulators) make responsibility easy to diffuse unless governance is explicitly designed.

Areas for improvement to explore with participants:

- balancing time between conceptual sessions and hands-on work,
- ensuring a shared baseline vocabulary early (particularly for mixed cohorts).

Next edition. A subsequent Winter School edition is planned for **January 2027**, and the feedback cycle will directly inform adjustments to curriculum sequencing, hackathon scaffolding, and interdisciplinary onboarding.

7) Conclusion

The Winter School on AI and Justice demonstrated that effective governance of AI systems requires more than technical optimization or formal legal compliance taken in isolation. By combining theoretical reflection, institutional perspectives, and a high-pressure applied hackathon, the programme helped participants practice a key lesson: **AI systems shape justice and the law not only through outputs, but through documentation, explanation pathways, contestability, and the distribution of responsibility.**

Next steps include:

- completing and synthesizing participant feedback,
- integrating findings into the design of the **January 2027** edition (especially around interdisciplinary onboarding and hackathon deliverables), and
- continuing to develop AI-NODES as a platform for research and training on the **normative and distributive effects** of AI in real-world decision infrastructures.

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

Gallery





Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

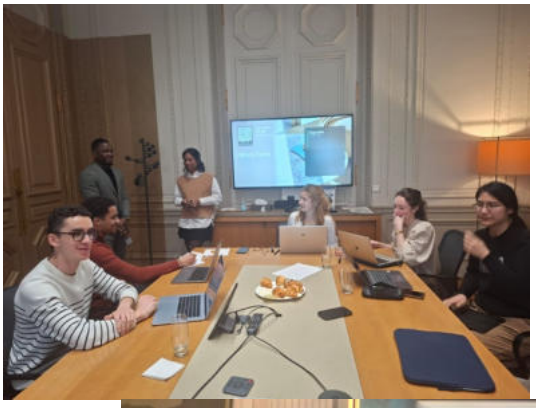

Testimonials from members of the winning team

Salim Cochereau

“As a first-year student in the joint master’s program in Law and Finance, participating in the Winter School on AI and Justice was an enriching experience. This one-week program was an introduction to one of the most emerging issues of our time: the transformation of law and justice in the age of AI.

Ce travail a bénéficié d'une aide de l’État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

The Winter School was a unique mix of theoretical insights and practical applications. It was the ideal platform for students like me who are eager to explore the intersection of technology and legal frameworks. We had the chance to think of how AI is reshaping our world and to examine how legal norms can—and must—guide this transformation.
The highlight of this program was the hackathon organized at Clifford Chance. This hands-on experience allowed us to apply the knowledge we had gained during the first days. Working alongside peers from both legal and technical backgrounds, we tackled real-world problems, the case of a discriminative algorithm widely used by banks across Europe. The hackathon reinforced our understanding of the subject matter but also fostered work cooperation and cross-field interactions.
This Winter School inspired me to continue exploring this dynamic field. I am grateful to Sciences Po, Sorbonne University, Clifford Chance and Google for creating such an impactful program!"

Eve Alaia Le Riche

"The Winter School Program on Justice and AI is a brilliant experience! The panel of courses offers an immersion in the technicalities of AI and Justice. From a *sciencespiste* perspective, it was great to explore the smaller, more grounded, picture as well as the bigger picture which we are more used to. Whether you are a neophyte or a becoming-specialist there will be material for you. Most importantly, the program is extremely fun. The week is tailored for us to enjoy our time. The Hackathon organised at Clifford Chance is an amazing experience. It was super exciting to play around with my team with everything we had learned. Finally, the visit at Google ended the week on an inspiring note. I was particularly touched by the discovery of the Google Arts and Culture Lab.

From this week, I gained motivation, much curiosity, and new friends!

Thank you to Sciences Po, Sorbonne University, Clifford Chance and Google for organising such a great program."

Ce travail a bénéficié d'une aide de l'État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.

Cécile Freiin von Roenne

“I took part in the Winter School on AI and Justice, jointly organised by Sciences Po and the Sorbonne Cluster for Artificial Intelligence, and it was a particularly valuable experience.

Across the week, we attended thought-provoking conferences on data, on how fairness is defined and operationalised in algorithms and AI systems, and on real-life case studies illustrating the concrete consequences of automated decision-making, especially in contexts involving fraud detection and discrimination.

We then put these concepts into practice during a hackathon hosted at Clifford Chance, where we acted as a regulating authority. This role forced us to take a balanced view, as we attempted to understand the perspectives of start-ups, banks, and affected individuals, and tested how regulatory decisions can remain both effective and fair when incentives and constraints diverge.

The programme also offered an in-house perspective through a visit to the Google France headquarters, including insightful sessions on copyright law and AI, which helped connect legal theory to how these issues are handled operationally inside a major technology company.

Many thanks to Rebecca Mignot-Mahdavi, Arnaud Latil, Séverine Dusollier, and all the organisers, Dessislava Savova and her team for hosting us at Clifford Chance, as well as Victor Fabre and his team at Google, for making the Winter School such a successful and insightful experience !”

Ce travail a bénéficié d'une aide de l’État gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0007.