

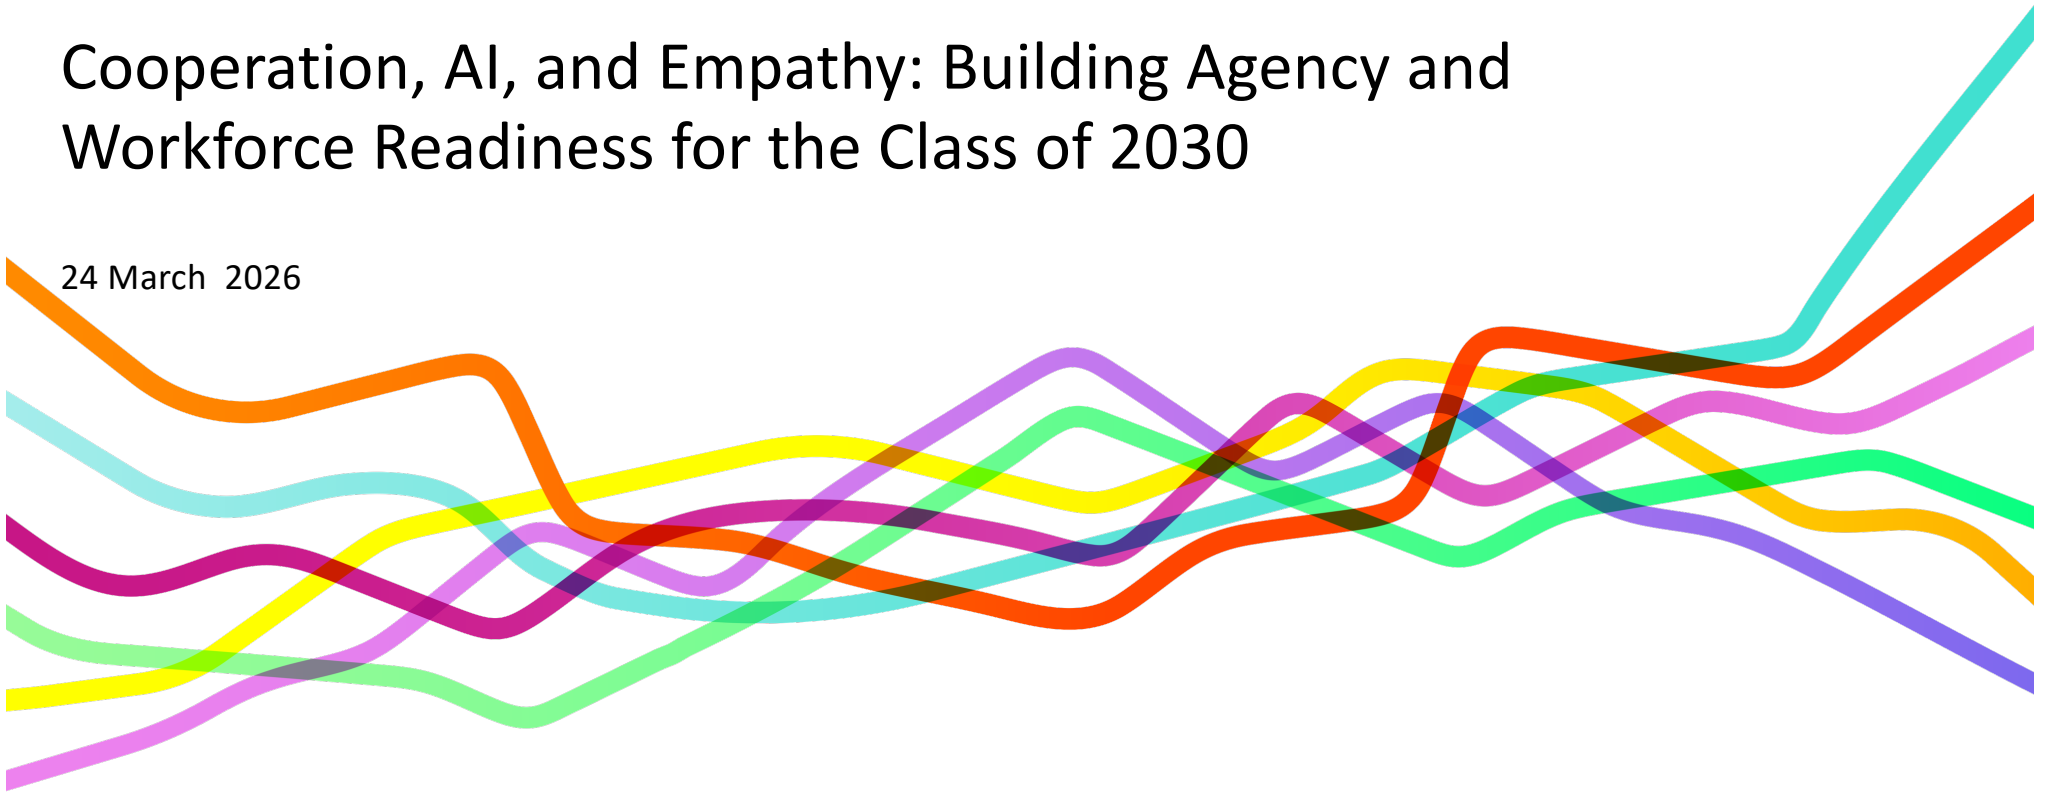
IMPERIAL

Centre for Languages,
Culture and Communication

Centre for Languages, Culture and Communication

Cooperation, AI, and Empathy: Building Agency and
Workforce Readiness for the Class of 2030

24 March 2026



IMPERIAL

AI and the Humanities: Designing Empathy into Digital Interactions

Situated within the after:hours programme, which brings together internal Imperial and external learners, this session links AI literacy to agency, employability, and professional identity.

Dr Tatia Codreanu

**CLCC after:hours, Institute of Lifelong Learning,
I-X AI in Science affiliate**

The Class of 2030

New work environments

The Class of 2030 will work in environments shaped by

- AI
- Automation
- Collaboration
- Rapid reskilling

Workforce readiness includes

- Technical fluency
- Empathy
- Service orientation
- The ability to use AI responsibly.

The Class of 2030

New work environments

The Class of 2030 will work in environments shaped by AI, automation, collaboration, and rapid reskilling. In this context, workforce readiness also depends on empathy, service orientation, judgment, and the ability to use AI responsibly.

AI readiness without empathy is incomplete. Productivity without agency is fragile.

Can students use AI in ways that strengthen agency, trust, and human cooperation?

The Class of 2030

Situated within the **after:hours** programme, CLCC

The **after:hours** programme creates a valuable bridge between higher education and professional life. It provides a shared space in which undergraduate and postgraduate students develop classroom learning into reflective, interdisciplinary, and work-facing practice. Working alongside adult and executive learners on the same course, they benefit from mutual exchange, broader perspectives, and cross-generational learning.

The course links AI use to placements, mentoring, employer expectations, and the articulation of value in GenAI-enabled workplaces.

The Class of 2030

From classroom learning to placements, mentoring, and enterprise

The course asks how interaction choices, such as tone, turn-taking, context-sharing, reciprocity, and clarity, shape outcomes in teamwork, leadership, and customer or patient-facing scenarios.

Can students use AI in ways that strengthen agency, trust, and human cooperation?

Students need more calibrated cooperation with AI. They need to know when AI is supporting good judgment, and when it is encouraging over-attribution, passivity, or avoidable error.

IMPERIAL

PART 1: The new reality, an AI economy built on unstable ground

From the illusion of collaboration to real cooperation

Evolutionary dynamics

GenAI often creates the experience of collaboration.

However, the system does not bear responsibility, liability, or downstream consequences.

This creates an illusion of cooperation.

Stable human-AI cooperation requires

- a User
- AI
- institution structure, with rules, monitoring, repair pathways, and accountability.

Case Study 1: AI & user as structural defectors

Mata v. Avianca, Inc. (2023)

THE EXPERIENCE

The AI was fluent, responsive, and seemingly cooperative. The lawyer experienced the interaction as authoritative legal guidance.

THE DEFECTION

The LLM hallucinated case citations, generating non-existent precedents that were filed in federal court.

THE CONSEQUENCE

Professional sanctions, judicial censure, and reputational damage absorbed entirely by the human cooperators, not the AI system. The judge sanctioned the lawyers and their firm jointly.

SOLUTION: build mandatory verification logs, audit trails, and shared organizational liability into deployment regimes.

Case Study 1 (contd.)

THE CORE PROBLEM

The Illusion of Cooperation

1

Generative AI systems operate with exceptional fluency in social signalling, triggering systematic overattribution of competence and authority.

2

Users experience responsiveness and apparent alignment, but the structural conditions for accountable cooperation are absent.

3

When errors occur, the asymmetry becomes visible: the model bears no liability. Human users absorb all consequences.

The AI had no skin in the game. It could defect without facing any consequences, while the human cooperators and the firm bear the full brunt.

The The Key Tension


PREDICTIVE COOPERATOR

AI is fluent, responsive, and seemingly aligned, performing cooperation without bearing any of its consequences.

RESPONSIBLE COOPERATOR

The human user experiences cooperation, verifies outputs, and absorbs all legal, professional, and reputational liability.

More details here →



Beyond the illusion of cooperation

TAXONOMY OF COOPERATION ROLES

Four actors shape whether cooperation succeeds or fails (Codreanu, 2026)

Responsible Cooperators

Humans who verify outputs, remain accountable, and maintain repair mechanisms. They actively increase governance and reduce liability.

Predictive Cooperators

AI systems generating contextually useful outputs aligned with user intent, but without bearing responsibility.

Over-Attributors

Users who treat fluent AI output as authoritative beyond its reliability. Reduce verification and increase uncorrected errors.

Structural Defectors

Organizations that deploy AI without adequate governance infrastructure.

<https://arxiv.org/abs/2602.19629>

What cooperative AI interaction looks like in practice

Interaction design

Polite AI is about designing interaction around:

- tone
- turn-taking
- context-sharing
- reciprocity
- clarity
- repair

Polite AI = empathy + formality → human-machine communication, with clear effects on trust, satisfaction, and adoption.

In many commercial AI systems, politeness is tuned for low-friction adoption: it often appears as reassurance, excessive politeness markers, and reduced challenge.

Design principle: empathy should support cooperation.

What cooperative AI interaction looks like in practice

Politeness alone is not enough

Excessive politeness can weaken trust if it becomes evasive.

Users need clarity over comfort in high-stakes or knowledge-driven tasks.

A better goal is **cooperative clarity** : Warm, respectful, honest, accountable interaction.

The risk is that AI performs cooperation at the level of style while weakening cooperation at the level of judgment.

The point is helping users think and act well.

Designing for agency

Autonomy, competence, relatedness

If we want students to become capable AI users, the interaction should support:

- Autonomy : the user feels in control.
- Competence: the user feels effective.
- Relatedness: the user feels socially connected and respected.

An agent that listens, affirms, and invites choices supports intrinsic motivation.

An agent that interrupts, assumes, or ignores undermines it.

This gives us a practical educational bridge from empathy to agency.

From teamwork to workforce readiness

The Cybernetic Teammate- field experiment with 776 professionals at Procter & Gamble ([Dell'Acqua et. Al 2025](#))

Key findings

- Individuals with AI matched the performance of human teams.
- AI-enabled professionals produced more balanced solutions across technical and commercial boundaries.
- Users also reported more positive emotion and motivation.

In terms of employability skills, tone, clarity, and emotional resonance shape human-AI teamwork, just as they shape human teamwork.

Cooperation theory

What defector-mode AI use looks like

- vague prompts
- no audience awareness
- no constraints
- no checking
- speed over sense-making
- over-attribution to fluent output
- organisational use without governance, audit trails, or escalation pathways

Likely outcome

- generic output
- misplaced trust
- poor transfer to real work
- weakened accountability
- avoidable error

Students need more calibrated cooperation with AI.

What my students learn instead

A simple cooperative prompt redesign habit

- Define the audience and context
- State the goal and stakes
- Ask for clarity (AI's prosocial alignment = adoption-oriented politeness)
- Invite alternatives or critique
- Check tone, inclusion, and accuracy
- Retain responsibility for the final output

Example class activity: *Draft a response to a frustrated stakeholder (patient, client etc.). Be clear, respectful, concise, and honest about uncertainty. Offer one next step. Then explain which parts of the response support trust and cooperation.*

Cultural fit and inclusive design

Politeness is not universal

English-centric or dominant-culture politeness norms can marginalise other users.
Some languages do not map neatly onto standard politeness markers.
AI can also reproduce gendered or class-based communication stereotypes.

Good design therefore needs

- culturally rich test scenarios
- diverse tone preferences
- human annotation
- transparency about trade-offs
- involvement from local communities and linguists where possible

A practical evaluation framework

What should we evaluate?

Did the system choose an appropriate tone for the context?

Did it use empathy, acknowledgment, or hedging appropriately?

Did it maintain clarity?

Did it reach a culturally and ethically acceptable outcome?

Did it support the user's agency?

A practical evaluation framework

Useful methods from our course materials

Code-based checks

LLM-as-a-judge

Human annotation

Politeness checklist : 1 to 5 rubric for clarity, empathy, courteousness, tone responsiveness, and user satisfaction.

Extending this into placements, mentoring, and enterprise

This framework fits naturally into

- placements
- mentoring
- employer briefs
- reflective portfolios
- careers guidance

Students can keep short AI reflection logs.

Mentors can ask how AI outputs were checked, adapted, and justified.

Employer-linked projects can assess value articulation, audience awareness, and responsible use.

This moves AI use from convenience to professional identity formation.

Building trust infrastructure around student AI use

Trust is an institutional outcome.


If we want students to become workforce-ready AI users, we should normalise

- reciprocity contracts
- visible traceability
- escalation and repair pathways
- training on limits and failure modes
- accountability leads in higher-stakes contexts

This can be operationalised through **three practical artefacts**

1. Human-AI Cooperation Charter
2. Defection Risk Register
3. Cooperation Readiness Audit

More details here →



The future-ready graduate

The person who can use AI with clarity, empathy, accountability, and judgment. Cooperation is the missing bridge between AI capability and workforce readiness. If we teach students to move from over-attribution and defector-mode habits toward reflective, responsible cooperation, we help them enter the labour market with stronger agency and stronger human skills.

Empathy is how AI readiness becomes humanly and professionally credible.

IMPERIAL

PART 2: Enterprise proof

Example of brands we study that operationalised these ideas in different ways. One protects identity and restraint. The other supports guided discovery and confidence.

How do we know an AI system improved human experience, judgment, and agency?

Case Study 2: AI as identity-sensitive cooperation

Brunello Cucinelli (Luxury Fashion) • Launched July 2024 •
brunellocucinelli.ai • 10,000–12,000 queries / day

Business problem: Luxury brands fear AI chat will dilute brand identity, a generic chatbot cannot carry the voice, philosophy or restraint that defines an ultra-premium brand.

THE CHALLENGE

AI risks diluting what makes a luxury brand feel luxury.

- Generic AI gives fast answers. Cucinelli's value is slowness, craft and philosophy, a standard chatbot is structurally incompatible with that identity.
- Internal teams were initially wary of AI. The deployment had to win over the organization.

THE AI APPROACH

A website without pages

- Solomei AI, bespoke platform built over 3 years by 2 mathematicians, 1 engineer, 1 artist, 1 philosopher. No menus, no index. Content flows in response to visitor intent.
- Live at brunellocucinelli.ai, ask anything, get answers, images and philosophy in real time.

WHAT MADE IT WORK

Identity-first governance

- Named AI agents act as 'guardians' (the Dioscuri), blocking off-topic questions. The AI cannot give political views or trend forecasts. Brand tone is non-negotiable by design.

Philosophy as infrastructure

- The site does not sell products. Its purpose is brand storytelling, the AI exists to deepen relationship.

AI can mean disciplined restraint, not friendliness alone. Cucinelli spent three years encoding philosophy, restraint and tone into the model. The governance layer (the 'guardians') was built before the features. <https://www.brunellocucinelli.ai/>

Case Study 3: AI as guided empathetic discovery

Jo Malone London (Estee Lauder Companies) • Launched December 2025 • Powered by Google Gemini + Google Cloud Vertex AI

Business problem: Online fragrance shopping lacks sensory cues — first-time buyers struggle to choose confidently without in-store expert guidance.

THE CHALLENGE

Fragrance is a sensory decision.

- Customers rely on in-store experts to translate mood and memory into scent. Online, that expertise disappears, leaving first-time buyers unable to choose confidently.
- The gap was harming digital conversion and undermining the brand's premium service promise in the D2C channel.

THE AI APPROACH

Conversational scent discovery

- Customers describe what they want. Google Gemini interprets responses and maps them to Jo Malone's proprietary olfactory data attributes to generate bespoke recommendations.
- Starts with four guided prompts. Trained on in-store consultation logic.

WHAT MADE IT WORK

Proprietary domain data

- Jo Malone's olfactory attribute database. Brand expertise encoded into the retrieval layer.

Tone and trust by design

- Conversational, playful UX mirrors in-store warmth. Builds confidence before purchase.

Interaction feels informed, trustworthy, and humanly legible. Jo Malone's advantage was the proprietary olfactory data and brand tone layered on top. The model was infrastructure. The knowledge was the product. <https://www.jomalone.co.uk/ai-scent-advisor>

Measuring Cooperative AI Experience

The bridge between Enterprise and educational frameworks

1

What human frustration or delay are we actually reducing?

2

Where does the user need visibility, consent or override power?

3

Which moments require a human checkpoint because consequences are material?

4

How will we know AI improved experience, not only throughput?

5

What evidence, logging and disclosure will make this defensible?

The future-ready graduate is the person who can recognise when AI is helping human cooperation, when it is weakening it, and what to change.

IMPERIAL

Thank you

**Cooperation, AI, and Empathy: Building Agency and
Workforce Readiness for the Class of 2030
24/03/2026**