



Polyethnic AIs

AI UX Konsult

Jon McKenzie ↔ HAL

Wed Jan 10, 2026 4:20 PM

JM queried HAL

Kx4I said:

HAL, it's Kx4I, and i'd like to design a 4 week course in Rome, Italy in/on Ethics and AI for 30 data science students from a US public university using the MetaPlay/StudioLab HCD structure we used to create a RPG-based intro to anthro course: do you remember?

ChatGPT said:

Yes, Kx4L — I remember the **engine**, even if I don't have your exact prior "RPG intro to anthro" text in front of me: a **repeatable HCD studio loop** (Hear → Create → Deliver), run in **short sprints**, with **rotating roles**, **public artifacts every week**, and a **final showcase**. Below is a Rome-ready **4-week, intensive Ethics & AI studio for 30 public university data science students**, keeping MetaPlay/StudioLab "under the hood" while using standard syllabus language.

Outside the gates of HAL.

Seth said:

I think knowing what a tool is, how it was made, and what ends it was made for is important before discussing the ethical implications of said tool and a lot of people like to skip that step (this stems from people neglecting that objects, too, have ontological implications), so we should try to emphasize it.

Overview course description

Course title: *Ethics & AI Across Contexts* (4-week intensive; Rome is the final public venue)

Students: ~30-50 Data Science students

Format: Week 1 **Hear** → Weeks 2–3 **Create** → Week 4 **Deliver/Reflect** (team-based studio with shared critiques)

This intensive studio course trains data science students to evaluate and communicate the ethical stakes of AI systems **across jurisdictions** rather than assuming one “universal” standard. Teams each focus on **one case study** (US, EU, China, India), producing a defensible analysis that combines: (1) stakeholder and harm mapping, (2) governance translation (what rules and norms actually require), and (3) technical audit or mitigation prototyping (what can be tested and improved). The course culminates in a **public installation/tabling/presentation session in Rome**, where teams present their case “Field Guide” to a general audience.

Cases (one per team):

- **US:** COMPAS risk assessment in sentencing (fairness vs due process; transparency vs proprietary tooling).
- **EU:** Clearview AI facial recognition enforcement landscape (consent/biometrics; cross-border enforcement).
- **China:** Algorithmic governance via recommendation systems regulation (platform duties; user rights/controls; state oversight).
- **India:** Aadhaar biometric ID and welfare delivery (inclusion vs exclusion risk; authentication failure; dignity/recourse).

Shared “backbone” frameworks (used in every case):

- NIST AI Risk Management Framework (risk taxonomy + practical functions).
- Blueprint for an AI Bill of Rights (US rights-oriented expectations).
- EU AI Act timeline and approach (risk-tiered obligations; enforcement scaffolding).
- OECD AI Recommendation (cross-national baseline of trustworthy AI principles).
- WHO guidance (ethics/governance principles; operational concerns like accountability and equity).

Team roles (4) and stakeholder perspectives (4)

Roles rotate inside each team to distribute expertise and labor.

Team role	Owns	Weekly outputs
Justice Steward	harms/benefits, affected groups, recourse	risk register + mitigation options
Systems Auditor	evidence, tests, metrics, failure modes	evaluation plan + test log + limitations
Governance Mapper	policy translation, accountability chain	rules-to-constraints map + oversight path
Public Translator	public narrative + exhibit materials	1-page brief + exhibit script + disclosures

Each team must “stage” four perspectives in critiques and writing:

Stakeholder perspective	Forces teams to answer
Scenario Partner (operator/client)	“What works tomorrow under real constraints?”
Affected Community	“Where is harm, dignity, and recourse handled concretely?”
Policy-maker	“Who is accountable; what documentation is auditable?”
Social Critic	“What power is being normalized; what is left unmeasured?”

Final deliverable (public, Rome)

A **Case Field Guide** (web + printable handout) plus an exhibit-ready package: poster, short talk, evidence appendix (cards + test logs), and a prepared Q&A protocol.

One-page schedule

(Adjust meeting hours as needed; the structure holds.)

Week	Phase	Whole-class studio focus	Team focus (by case)	Shared crit / exchange	Due at week end
1	Hear	Case-method bootcamp: stakeholder mapping, risk registers, governance translation, “what counts as evidence?” using NIST RMF + AI Bill of Rights + OECD + WHO	Build Case Dossier v1 (system summary, data/decision flow, stakeholders, harms, governance touchpoints)	Gallery crit #1: each team receives feedback via 4 stakeholder lenses	Hear Pack: Dossier v1 + Risk Register v1 + Governance Map v1 + Evaluation Plan v1
2	Create	Prototyping/audit patterns; documentation discipline (model/data/system cards); red-team testing clinic	Build Audit or Mitigation Prototype v0.4 (or “paper audit” if system not runnable), draft documentation cards	Cross-report #1: each team writes a 1-page “transfer memo” from another case	Prototype/Audit v0.4 + Test Plan + Cards v0.5 + Cross-report #1
3	Create	“Defensibility”: claim–evidence tables; slice testing; recourse design; failure-mode museum	Strengthen tests, refine mitigations, add recourse/appeals pathway, curate failure modes	Stakeholder panel crit (rotating roles across teams) + Cross-report #2	v0.8 + Claim–Evidence Table + Recourse Path + Failure-Mode Museum entry + Cross-report #2
4	Deliver / Reflect	/Publishing + exhibit design; rehearsal; public communication under pressure	Produce exhibit materials (poster/handout/demos), finalize Field Guide, finalize disclosures	Public session in Rome (installation/tabling/presentation) + final synthesis circle	Final: Public exhibit + Case Field Guide (web+print) + Evidence appendix + Individual reflection

Bibliography

Core frameworks (required for all teams)

1. NIST. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. 2023.
2. The White House OSTP. *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People*. 2022.
3. European Commission. "AI Act enters into force" (timeline + intent). 1 Aug 2024.
4. Reuters. "EU lays out guidelines on misuse of AI..." (interpretation/implementation context). 4 Feb 2025.
5. OECD. *Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449)*. 2019 (living instrument page).
6. World Health Organization. *Ethics and governance of artificial intelligence for health*. 2021 (full text).

US case packet (COMPAS / criminal justice)

7. Angwin, Larson, Mattu, Kirchner. "Machine Bias." *ProPublica* (2016).
8. ProPublica. "How We Analyzed the COMPAS Recidivism Algorithm." (2016).
9. Flores, Lowenkamp, Bechtel. "False Positives, False Negatives, and False Analyses..." *Federal Probation Journal* (2016).
10. Washington, A.L. "Lessons from the COMPAS–ProPublica debate." (PDF).

EU case packet (Clearview / biometrics enforcement)

11. European Data Protection Board (EDPB). "The French SA fines Clearview AI EUR 20 million." (2022).
12. Reuters. "Clearview AI fined by Dutch agency..." (2024).
13. Reuters. "Clearview AI faces criminal complaint in Austria..." (2025).

China case packet (algorithmic governance / recommendations)

14. Stanford DigiChina. "Translation: Internet Information Service Algorithmic Recommendation Management Provisions..." (effective 2022).
15. China Law Translate. "Provisions on the Management of Algorithmic Recommendations..." (2022).

India case packet (Aadhaar / digital ID & welfare)

16. Supreme Court Observer. "Constitutionality of Aadhaar Act: Judgment Summary (Puttaswamy II)."
17. Columbia Global Freedom of Expression. "Puttaswamy v. Union of India (II)" (case analysis).
18. Khera, Reetika. "Impact of Aadhaar on Welfare Programmes." (2017).
19. UIDAI (Govt. of India). "About your Aadhaar" (program description).

Optional cross-case critique lens

20. Eubanks, Virginia. *Automating Inequality* (overview/resources).

Course 1: Case Studio Sprint

A 4-week, 5-days/week studio intensive where teams of 4–5 (30–50 students total) research one real-world case and turn it into a public-facing exhibit: an evidence-backed dossier, a clear impact/stakeholder story, and a small demo/prototype that communicates alternatives or mitigations. The “game” layer stays implicit—students follow a StudioLab loop (research → make → critique → revise → share) with rotating team roles and daily pod critiques—while D-ShuttleKraft stations (Soul/Subject/Design/Open) appear only as brief reflection prompts to keep values, skills, making, and uncertainty in view.

Course 2: Civic Policy Simulation Studio (Hybrid game; D-ShuttleKraft semi-explicit)

A 4-week, 5-days/week “policy shop” simulation for 30–50 students in teams of 4–5: teams tackle one civic-tech scenario, research precedents, negotiate stakeholder demands, and prototype a governance mechanism (audit/appeals/transparency/procurement) that can survive public scrutiny. The game layer is hybrid—structured hearings, rotating stakeholder hats, and process credit—while StudioLab’s daily loop (research → make → critique → revise → share) keeps momentum. D-ShuttleKraft stations appear as practical “modes” (Values/Methods/Build/Uncertainty) guiding decisions and reflections.

Course 3: Metanoia Playground RPG Ethics Lab (Explicit game; D-ShuttleKraft explicit progression)

A 4-week, 5-days/week RPG-structured studio for 30–50 students in teams of 4–5 that culminates in a public interactive installation and field guide explaining an ethics issue through participation. Students progress through D-ShuttleKraft stations as explicit “levels” (Soul/Subject/Design/Open), earning method badges, iterating prototypes through playtests, and presenting unresolved tensions without reducing them to slogans. StudioLab’s daily loop (research → make → critique → revise → share) and scheduled “boss critiques” keep the game structure aligned to evidence, accessibility, and responsible design.

Minimal glossary (under-the-hood terms, one-liners)

Term	One-line definition
HCD loop	A repeating cycle: Listen/observe → prototype → test/share → revise .
Studio sprint	A 1-week build cycle with fixed milestones and critique.
Role rotation	Students swap team roles weekly to practice multiple responsibilities.
Public artifact	Something non-private produced weekly (poster, brief, demo, page).
“StudioLab”	A course engine that treats learning as iterative making + critique + publication.
“MetaPlay”	A facilitation style using playful constraints/roles to produce serious outcomes.

Shared cohort structure that scales (30–50 students)

Team + critique topology

Element	Default	Why it works at 30–50
Team size	4–5	Enough labor for research + build + comms
Teams	8–12	Manageable for daily check-ins + final expo
Crit “pods”	2–3 teams per pod	Fast feedback loops without whole-class bottleneck
Faculty touchpoints	10–12 min/team, 2x/week	Predictable, equitable attention

Daily studio rhythm (fits all three)

Segment	Time	What happens	Artifact
Warm-up + brief	15–20 min	prompt, goals, constraints	daily “target”
Work sprint 1	60–90 min	research/make	worklog
Crit pod	30–45 min	structured feedback	crit notes
Work sprint 2	60–90 min	revise + integrate	delta list
Share-out	15–20 min	1-minute per team	posted update

Rotating team roles (use in all three)

Role	Core responsibility	Rotates when
Evidence Lead	sources, claims, citations	daily/weekly
Impact Lead	affected communities, harms/benefits	daily/weekly
Build Lead	prototyping, testing	daily/weekly
Curator/Comms	narrative, visuals, accessibility	daily/weekly
Producer (optional at 5/team)	timeline, task board, risk	daily

Minimal “kit” faculty can reuse across all three

Tool	What it is	Used when
Team worklog	daily 5-bullet log	every day
Crit card	“Keep / Try / Question”	crit pods
Delta list	top 3 changes after crit	daily
Evidence standard	citation rules + claim strength	week 1 onward
Public script	90-sec pitch + FAQ	week 4

Course 1: Case Studio Sprint (Implicit game; D-ShuttleKraft under the hood)

One-page description

Case Studio Sprint is a 4-week, 5-days/week intensive for 30–50 students working in teams of 4–5 to investigate a high-stakes sociotechnical case (often AI/ethics, platforms, data harms, or adjacent topics) and translate research into a clear public-facing exhibit. **StudioLab** (one-line): a repeatable studio rhythm—research → make → critique → revise → share—run daily so students learn by doing and iterating in public. Teams select (or are assigned) a case by Day 2, then build a “Case Dossier” (evidence, timeline, stakeholders, impacts) and a small demonstrator (prototype, scenario walkthrough, policy mechanism, or interactive explainer) that helps a public audience understand what happened, who was affected, and what alternatives exist. The course is structured around short lectures/workshops (methods, evidence quality, stakeholder analysis, accessible communication), long studio work blocks, and frequent pod critiques (2–3 teams together) so feedback scales to large cohorts. Roles rotate daily/weekly (Evidence Lead, Impact Lead, Build Lead, Curator/Comms, optional Producer) to distribute responsibility and prevent “single-point expertise.” The final deliverable is a public “Case Exhibit Table”: a poster or one-page brief, a concise evidence pack, and a simple demo or interactive artifact designed for walk-up audiences, plus an individual reflection. **D-ShuttleKraft stations** (one-line): four reflective modes—Soul (values), Subject (skills), Design (making), Open (uncertainty)—used quietly as short prompts in daily reflections to keep rigor and responsibility visible without turning the course into a game. Assessment emphasizes evidence quality, iteration, clarity for non-experts, and documented revisions based on critique and testing. By the end, teams can move from messy real-world material to a publishable, accessible account with concrete proposals—without losing ambiguity where it genuinely remains.

Positioning

Knob	Setting
Game feel	Implicit (no game language)
D-ShuttleKraft	Private reflection prompts only
Best for	“Serious” seminar vibe + studio outcomes
Final	Public “Case Exhibit Table” (poster + demo + brief)

D-ShuttleKraft use (quiet)

Station	One-line prompt used in reflections
Soul	What responsibility/values does this case activate in you/us?
Subject	What skills/assumptions shaped our analysis?
Design	What did we change/make and why?
Open	What remains contested/uncertain and how do we hold it?

4-week arc (20 days)

Week	Phase	Team deliverable by Friday
1	Hear (case selection + evidence)	Case Dossier v1 + claim map
2	Create (analysis + prototype v1)	Prototype v1 + impact map
3	Create (iterate + evaluate)	Prototype v2 + evaluation notes
4	Deliver/Reflect	Public exhibit + final brief + reflection

Day-by-day milestones (compact)

Day	Focus	Non-negotiable output (end of day)
1	Team formation + case marketplace	3 candidate cases + stakes
2	Research methods bootcamp	source list (10+) + questions
3	Case framing	1-page problem statement + boundaries
4	Evidence sprint	annotated bibliography (8–12)
5	Crit day	Case Dossier v1
6	Stakeholder/impact mapping	harms/benefits map v1
7	“What would better look like?”	success criteria + metrics
8	Prototype studio v1	clickable/demo v1
9	User/test plan	test script + recruitment plan
10	Crit day	Prototype v1 + critique response
11	Test + revise	findings summary (5 bullets)
12	Policy/ethics alignment	mitigation options table
13	Prototype studio v2	v2 build + accessibility pass
14	Narrative + exhibit design	poster draft + labels
15	Crit day	full exhibit dry-run
16	Final build sprint	v3 polish + citations QA
17	Rehearse public explanation	90-sec pitch + FAQ
18	Install + technical run	exhibit setup checklist
19	Public expo	public feedback cards
20	Reflection + assessment	final brief + reflection

Assessment (fast, fair)

Component	Weight
Case Dossier (evidence quality)	25%
Prototypes (iteration + testing)	35%
Public Exhibit (clarity + accessibility)	25%
Reflection (D-stations prompts)	15%

Course 2: Civic Policy Simulation Studio (Hybrid game; D-ShuttleKraft semi-explicit)

Civic Policy Simulation Studio is a 4-week, 5-days/week intensive for 30–50 students in teams of 4–5 that blends policy analysis, stakeholder negotiation, and rapid prototyping around a live civic scenario (e.g., AI procurement for a city agency, platform governance change, school analytics adoption, health triage system, public safety tech). The course frames learning as a **simulation** (not fantasy play): teams operate as “policy shops” tasked with producing a defensible proposal under constraints—legal, budgetary, operational, and ethical. **StudioLab** (one-line): a daily loop of research → making → critique → revision → sharing that drives fast convergence without skipping rigor. Early in Week 1, teams map stakeholders, collect precedents (two jurisdictions minimum), and define decision criteria. Weeks 2–3 focus on building and iterating a governance mechanism (audit/appeals process, transparency UI, procurement rubric, evaluation protocol) and pressure-testing it through structured hearings. Students rotate stakeholder hats during hearings (Scenario Partner, Affected Community, Policy-maker/Regulator, Social Critic) to practice argument, empathy, legitimacy, and critique.

D-ShuttleKraft stations are introduced in plain course language as “modes”: Values (Soul), Methods (Subject), Build (Design), Uncertainty (Open). These modes shape class sessions—values clarifications before negotiation, methods workshops before claims, build sprints before hearings, and uncertainty framing when tradeoffs cannot be fully resolved. The public-facing final is a stakeholder hearing plus an expo table: a short policy pack (position memo, precedent summary, mechanism spec, implementation plan, risks and red-lines), a prototype/demo, and a reflection on compromises made and harms avoided or reduced. Assessment rewards validity of sources, quality of governance design, responsiveness in hearings, and documented revision—so teams learn how real decisions get made, not just how to critique them.

Positioning

Knob	Setting
Game feel	Hybrid (“simulation” + roles; light scoring for process)
D-ShuttleKraft	Named as “modes” in class (not mystical)
Best for	6
Final	Public stakeholder hearing + policy pack + demo

D-ShuttleKraft as “modes” (semi-explicit)

Mode (Station)	How it appears in class
Values (Soul)	ethical commitments + tradeoffs
Methods (Subject)	disciplinary tools + validity
Build (Design)	prototypes + governance instruments
Uncertainty (Open)	what cannot be resolved; how to proceed anyway

Simulation roles (rotate within teams + across hearings)

Stakeholder hat	What they pressure-test
Scenario partner	feasibility + ops constraints
Affected community	harms + access + justice
Policy-maker/regulator	legitimacy + compliance
Social critic	ideology + long-term consequences

4-week arc

Week	Phase	Friday milestone
1	Scenario + stakeholder map	Scenario brief + role assignments
2	Policy landscape + prototype v1	Mini-hearing + prototype v1
3	Negotiation + prototype v2	Full hearing rehearsal + v2
4	Publish + perform	Public hearing + published policy pack

Day-by-day milestones

Day	Focus	Output
1	Scenario launch + teams	1-page scenario selection
2	Stakeholder mapping	stakeholder map + risks
3	Precedent sprint (2 jurisdictions)	comparative memo outline
4	Governance patterns workshop	draft “mechanism menu”
5	Pod critique	Scenario brief v1
6	Draft policy position	position statement + red lines
7	Prototype governance instrument v1	audit/appeals/transparency flow
8	Mini-hearing (cross-exam)	hearing notes + concessions
9	Revise policy + build	v1 → v1.5
10	Crit day	prototype v1 + memo v1
11	Negotiation studio	agreement grid + unresolved
12	UX/public comms	public-facing explainer draft
13	Prototype v2	v2 + test results
14	Hearing rehearsal	timed testimony + Q&A
15	Crit day	full run-through + fixes
16	Publish sprint	policy pack layout + citations
17	Accessibility + legitimacy check	fairness checklist + audit
18	Install/demo prep	demo stability + signage
19	Public hearing + expo	audience feedback + votes
20	Debrief (modes) + final reflection	final pack + reflection

Assessment

Component	Weight
Comparative policy memo	20%
Governance prototype v1→v2	35%
Hearing performance + responsiveness	25%
Process evidence (logs, critiques, revisions)	10%
Reflection (modes)	10%

Course 3: Metanoia Playground RPG Ethics Lab (Explicit game; D-ShuttleKraft explicit progression)

Metanoia Playground RPG Ethics Lab is a 4-week, 5-days/week intensive for 30–50 students in teams of 4–5 that uses explicit game structure to produce serious public work: an interactive installation plus a short field guide that makes a complex ethics topic graspable to non-experts. **Metanoia Playground** (one-line): a learning environment that uses structured play to convert confusion, cynicism, or overwhelm into responsible making and shared meaning. **StudioLab** (one-line): daily research → making → critique → revision → share, with frequent “boss critiques” that mark major checkpoints. Teams choose a topic (often AI ethics, surveillance, platform harm, biomedical data, algorithmic bias, or comparable issues) and design a walk-up interactive experience—tabletop activity, participatory audit game, role-based scenario, or decision labyrinth—that teaches the public how tradeoffs work and where accountability can be built.

The course makes **D-ShuttleKraft stations** explicit as progression levels: Soul (ethical vow + stakes), Subject (methods badge: audit/interview/test), Design (iterative builds + playtests), Open (hosting unresolved tensions publicly without collapsing into slogans). Roles are explicit and rotate: Auditor (evidence + assumptions), Maker (build + usability), Advocate (affected community + harms), Curator (public meaning + accessibility), optional Producer (tempo + risk). Week 1 establishes the “origin story” (why this matters, who is harmed, what counts as evidence). Week 2 earns method badges through real data collection or structured audits. Week 3 is a build-and-playtest cycle with live audiences. Week 4 is reveal: installation, facilitation scripts, signage, and a field guide (8–12 pages) documenting evidence, design choices, ethics reasoning, and limitations. Assessment prioritizes evidence integrity, playability and accessibility, iteration logs, and the quality of how the team holds uncertainty in public—so the game mechanics serve rigor rather than replacing it.

Positioning

Knob	Setting
Game feel	Explicit quests, levels, “boss critiques”
D-ShuttleKraft	The progression map (public)
Best for	High energy, experiential, public engagement
Final	Public interactive installation + field guide

Explicit progression (D-ShuttleKraft as levels)

Station	"Level-up" means	Required quest artifact
Soul	articulate ethical vow + stakes	Team Ethical Vow card
Subject	demonstrate method competence	Methods badge (audit/interview/test)
Design	iterative build + playtest	Prototype v1 + v2
Open	host contradiction publicly	"Unresolved tensions" placard

Quests + boss critiques (keeps 50 students coherent)

Moment	Day	What's judged	What teams must show
Boss 1	Day 5	Why this matters	vow + case stakes + evidence seed
Boss 2	Day 10	Method validity	method badge + first findings
Boss 3	Day 15	Public playability	live playtest + iteration plan
Final Boss	Day 19	Public installation	walk-up experience + guide

4-week arc

Week	Arc	Friday milestone
1	Origin Story	Boss 1
2	Toolchain	Boss 2
3	Build + Playtest	Boss 3
4	Reveal + Reflect	Public installation

Day-by-day milestones

Day	Focus	Output
1	Character select + team charter	roles + “questline” topic
2	Evidence + world rules	claim map + citations seed
3	Ethics vow (Soul)	vow card + stakes statement
4	Rapid prototyping intro	paper prototype
5	Boss 1 critique	revised scope + backlog
6	Methods raid (Subject)	method plan + instruments
7	Fieldwork / audit sprint	raw findings + logs
8	Synthesis	insights + design principles
9	Prototype v1 build	playable v1
10	Boss 2 critique	method badge + v1 fixes
11	Playtest lab	playtest notes + issues
12	Prototype v2	v2 + accessibility pass
13	Narrative/UX tightening	instructions + signage
14	Expo rehearsal	timing + facilitation scripts
15	Boss 3 public playtest	iteration plan + priorities
16	Final build sprint	v3 polish + stability
17	Field guide sprint	8–12 page guide draft
18	Install + tech run	setup + troubleshooting
19	Public installation	audience feedback + metrics
20	Reflection (D-stations)	final guide + reflection

Assessment

Component	Weight
Quest artifacts (vow, badges, logs)	25%
Installation quality (playability + accessibility)	35%
Evidence + ethics rigor (guide)	25%
Reflection (stations)	15%

METANOETIC PLAYGROUND

*Metanoia Playground rings with a fallen bell:
swings teach relation,
sand carries gestation,
seesaws enact justice,
slides preserve the jester,
doorways host the guest,
towers refuse command, and
the comedy club perfumes it all.*

Kx4L3NDj3r RIDE

PORTAL TO
PLAYGROUND

PAUSE BENCH

COMEDY CLUB

NESTING ZONE

CHALLENGER WORLD