

[YOU] on AI + X

Berkeley Cohort - AI Mindset, New Human Skills, and an AI Project

Instructor of Record at UC Berkeley: Trond Petersen, Sociology, UC Berkeley

with the Instructional Team: Alex Bayen (Berkeley EECS),

Edo Segal (CTO Napster), and Ikhlaz Sidhu (Dean IE University, Senior Fellow CITRIS/Berkeley)

This course, offered for research credit, is a project-and-lecture experience that culminates in both project delivery and participatory homework completion. Students who participate receive course or project credit from their home institutions:

- **At UC Berkeley:**
Sign up for SOCIOL 198.11 for undergraduates. 2 units.
Sign up for SOCIOL 298.1 for graduate students. 2 units.
- **Parallel Course at IE University:** 3 ECTS credits from the School of Science & Technology, through the IE Silicon Valley Gateway Program. Lead Instructor Ikhlaz Sidhu, with teaching team Trond Petersen, Alex Bayen and Edo Segal. Contact the program management office at IE Sci-Tech.

Regardless of home institution, students must first be selected for the IE Silicon Valley Hub program through an application process. The application is available at: <https://yaix.netlify.app/>

Course Title	[YOU] on AI + X — AI Mindset, New Human Skills, and an AI Project
Instructors	Instructor on record: Trond Petersen, with guest instructors Edo Segal (ES), Alex Bayen (AB), Ikhlaz Sidhu (IS), and coach)
Units / Credit	UC Berkeley: 2 units of credit, sponsored by Prof. Trond Petersen (L&S) and Prof. Alex Bayen (Engineering).
Level / Eligibility	Open to all majors.
Term	[Term, Year]
Format & Duration	8 instructional modules plus a pre-course orientation (Module 0). Each module comprises two 45-minute sessions and one 90-minute project/lab session.
Meeting Days / Times	Mondays, 9:30–11:00 AM PT / 6:30–8:00 PM CET — Sessions 1 and 2 Wednesdays, 9:30–11:00 AM PT / 6:30–8:00 PM CET — Session 3 and optional lab time
Modality / Location	Synchronous sessions held jointly across Pacific and Central European time zones; room, platform, and modality (in person, hybrid) confirmed per host site. [TBD]
Course Platform	The [YOU] on AI companion platform — an AI-agentic workspace including the Edo and Ikhlaz digital-twin mentors, AI-led assignments customized per student, and the [THEM] on AI library and AI Field Guide.
Grading Basis	Pass / No Pass (research credit)

Note: bracketed items ([TBD], [Term, Year], [IE email]) are institutional fields set per offering. This syllabus is subject to change; students will be notified of any revisions in advance.

COURSE OVERVIEW

Course Description

This course is built on a premise it then asks students to test: that the most valuable contribution a university can make in the age of artificial intelligence is not information delivery — language models now provide that on demand — but formation: the cultivation of judgment, taste, durable human skills, and the ability to decide, and defend, what is worth building. The course opens with a foundational question — what are we doing here, and what will you accomplish this term, given the pace of change in AI? — and spends the term testing possible answers through direct, applied practice.

The course design draws on three sources: Edo Segal's account of artificial intelligence as an amplifier of human capacity, developed in his book [YOU] on AI; Ikhlāq Sidhu's framework for innovation mindset and skill-building, drawn from the Berkeley Method of Entrepreneurship and the Data-X curriculum; and a hands-on AI project that requires students to practice both. As in Data-X, the course is roughly half technical and half human/innovation in focus. At its core, it is an advanced, highly applied project course: students ideate a project in the first half of the term and iterate it to completion in the second.

The course also, in practice, engages in an active debate about the future of higher education in an AI-saturated environment — namely, whether the university's distinctive value lies in the deliberate cultivation of judgment and character rather than in the transmission of facts and techniques alone. Students should expect the course to model, rather than merely discuss, this position.

Course Design: Two Tracks Per Module

Each of the eight instructional modules pairs a conceptual strand with an applied strand: students read and discuss a worldview grounded in the assigned text, and in the same week practice the AI tool set and project work that puts the idea into use. Concretely, each module runs three parallel tracks:

- **Broader Insight — human skills & worldview:** a lecture built around the recurring vocabulary of the course (e.g., the river and the beaver, friction and flow, the software death cross, the candle, the silent middle) and two chapters of the assigned text, closing with a mindset objective and a short individual assignment.
- **Code & Theory — the AI toolset:** just-in-time technique in support of the student project, covering the data/AI pipeline, prompting and co-intelligence, agentic build, and evaluation.
- **Teams & Projects — the studio:** live team check-ins, structured critique, and the iterative development of a real AI artifact in partnership with an external mentor.

Foundational Mindset Shifts

Five recurring shifts run through the course and shape what “competence with AI” is taken to mean:

1. Intelligence is not confined to the individual; it is distributed across tools, agents, and networks — the practical skill is learning to think with AI, not to have AI think for you.
2. The goal of a project is not a single output but a repeatable system: “build the factory, not the shoe.”
3. A student's working network increasingly includes AI agents alongside people, and both must be managed deliberately.
4. As AI performs more routine tasks, the differentiating skill shifts from executing tasks to directing intelligence toward a well-chosen end.

5. AI amplifies whatever signal it is given—for better or worse—which places a premium on the quality of a student's questions, purpose, and judgment.

Underlying all five is a broader, empirically observable trend the course treats as a working assumption: skills increasingly outweigh pedigree in determining what a person can accomplish.

STUDENT LEARNING OUTCOMES

By the end of the course, students will be able to:

1. Articulate and defend a personal position on the role of human judgment in an AI-saturated environment, supported by evidence from the assigned text and from their own project work.
2. Design, build, and iterate a working AI-enabled artifact — a venture prototype, a research tool, or a creative or professional artifact — in a team, in partnership with an external mentor, from ideation through a live demonstration.
3. Apply core AI and data-pipeline techniques, including prompting and co-intelligence practice, agentic workflow design, and evaluation, selecting technology-independent approaches on the basis of evidence rather than default preference.
4. Demonstrate and reflect on growth in seven human and innovation skills — curiosity and motivation, critical thinking, grit and deliberate practice, emotional intelligence, comfort with discomfort, trust and psychological safety, and self-knowledge — as measured by validated pre/post instruments.
5. Distinguish AI-generated scaffolding from AI-generated substitution, and plausible output from verified output, applying a structured critical-thinking framework to identify hallucination and bias.
6. Conduct a stakeholder and dual-use risk assessment of an AI system, identifying potential benefits, harms, and the human safeguards required before the system scales.
7. Disclose and defend the role of AI assistance in their own work, maintaining authorship and accountability for all submitted material.

PROJECT TRACKS

Students choose one project track in Module 2 and carry it through to Demo Day in Module 8. All three tracks share a common arc: ideation and evidence-gathering in the first half of the course, and iterative building toward a live demonstration in the second.

Venture — Entrepreneurship On / With AI

Build an AI-enabled product or service for a real user or market. Deliverable: a working prototype together with an opportunity and business story (who it serves, and why it wins). Maps to lean-startup methodology and the Business Model Canvas.

Research — Discovery On / With AI

Use AI to investigate a real question: gather and analyze data, run a structured inquiry, and produce an evidence-backed finding. Deliverable: a working analysis or tool together with a research note stating methods, results, and honest limitations.

Studio — Creation and Craft On / With AI

Make something: a creative or professional artifact — media, design, writing, an application, or a public-good tool. Deliverable: the finished artifact together with a maker's statement on the human judgment behind it. Social-impact projects are welcome in this track.

THE HUMAN-SKILLS TRACK

Every module develops one human or innovation skill, taught through Berkeley-Method inductive exercises, engineered discomfort, and structured reflection, and measured with validated instruments. Progress on this track is graded on its own rubric; it is assessed in parallel with, not folded invisibly into, the project grade.

- **Curiosity & motivation** — grounded in Self-Determination Theory; question-generation exercises; articulating a personal “why.”
- **Critical thinking** — the Paul–Elder framework; detecting hallucination and bias; the principle that plausible is not true.
- **Grit & deliberate practice** — grounded in Duckworth's research; responding to engineered setbacks; measured with the Grit-S scale, pre- and post-course.
- **Emotional intelligence** — the RULER framework; emotion labeling; care for users and teammates; measured with an EQ instrument, pre- and post-course.
- **Comfort with discomfort** — engineered discomfort exercises (cold outreach, public pitch) paired with structured reflection.
- **Trust & psychological safety** — grounded in Edmondson's research; team retrospectives and structured peer feedback.
- **Self-knowledge** — measured with the Berkeley Innovation Index, administered pre- and post-course as a growth measure.

REQUIRED TEXTS AND RESOURCES

- *[YOU] on AI* by Edo Segal — the core text (full text available on the companion platform). Standard price: 40 USD, but the electronic version will be available at no cost during the trial course.
- Ikhlāq Sidhu — selected materials on Innovation Engineering and the Berkeley Method; Data-X resources.
- The [THEM] on AI library and the AI Field Guide (10,000+ entries), at youonai.ai.
- The course companion platform, including AI digital-twin mentors of Edo Segal and Ikhlāq Sidhu, and AI-led assignments customized to each student's project and growth.

COURSE PLATFORM AND COMMUNICATION

Announcements and team conversation take place on the companion platform and the class communication channel. A significant portion of student learning is expected to occur through peer interaction and through the AI twin-mentors; both are treated as part of the course method, not as supplements to it.

COURSE SCHEDULE

The schedule below outlines course material and structure and may shift as the term unfolds; students will be notified of any changes in advance. Each session includes a ten-minute break.

Schedule Overview

No.	Theme	Key Deliverable
0	Orientation — Introduction to [YOU] on AI + X	—
1	Formation Over Information — the Foundational Question	Team-formation note + 3 candidate problems
2	The Amplifier — Motivation, Voice, and Co-Intelligence	Project proposal (individual + team)
3	Judgment Under Evidence — Critical Thinking and the State of the Art	Problem statement + NABC story + SOTA scan
4	The Candle — Authorship, Judgment, and Discomfort	Working prototype (“factory, not shoe”)
5	Friction, Flow, and Scope Lock — From Exploration to Execution	Mid-course demo (v0) + build plan
6	Stewardship — Risk, Ethics, and Observability	Risk/ethics audit + progress demo
7	Attention and Care — Human-Centered Evaluation	Human-in-the-loop observation memo
8	Signing Your Work — Demo Day and the Foundational Question, Revisited	Demo Day presentation + final reflection

Detailed Weekly Schedule

Module 0 — Orientation — Introduction to [YOU] on AI + X

Broader Insight — Mindset & Worldview.

Lecture: A pre-course orientation to the premise, structure, and two-track (mindset / build) design of the course, delivered before Module 1 begins.

Reading: None assigned.

Assignment: None assigned.

Code & Theory — AI Tools & Practice.

Practice: First access to the course companion platform and the AI digital-twin mentors.

Assignment: None assigned.

Teams & Projects.

In Class: None — team formation begins in Module 1.

Assignment: None assigned.

Module 1 — Formation Over Information — the Foundational Question

Broader Insight — Mindset & Worldview.

Lecture: What are we doing here, and what will you accomplish this term, given the pace of change in AI? Introduces formation over information, the forge, the fishbowl, and AI as amplifier. Mindset objective: intelligence is not only inside you, it is around you — learn to think with AI, not have AI think for you.

Reading: [YOU] on AI, Ch. 1 “The Winter Something Changed” and Ch. 2 “The Discourse.”

Assignment: Understand You and AI — read the assigned chapters and write a first answer, in your own words, to the foundational question.

Code & Theory — AI Tools & Practice.

Practice: Meet the Edo and Ikhlāq AI digital twins; practice using AI as tutor, coach, teammate, tool, simulator, and student.

Assignment: Be Ready with AI Tools — set up access to at least one capable AI system or the companion platform and use it to support the week's work. Submit a short AI-use disclosure log identifying where AI helped you think, explore, organize, or critique, and where you kept final judgment.

Teams & Projects.

In Class: Form diverse teams; survey the Venture, Research, and Studio tracks; bring three problems you care about and map them to possible users.

Assignment: Team-formation note plus three candidate problems per student.

Module 2 — The Amplifier — Motivation, Voice, and Co-Intelligence

Broader Insight — Mindset & Worldview.

Lecture: The inversion and the amplifier: the machine has learned language, so signal quality now matters more. Motivation and curiosity as the engine of work; voice, change, and agency as the posture for building. Mindset objective: AI amplifies you, for good and bad, so the quality of your questions, purpose, and judgment matters.

Reading: Ch. 3 “When the Machine Learned Our Language” and Ch. 4 “Dylan's Like a Rolling Stone.”

Assignment: Motivation — write a personal motivation note on why the chosen problem matters to you.

Code & Theory — AI Tools & Practice.

Practice: Co-intelligence in practice: prompting as thinking, the jagged frontier, and using AI to widen, challenge, and critique rather than simply answer.

Assignment: Think with AI — use AI to map the problem space, then add your own critique of its blind spots; iterate toward a fuller picture of users, causes, stakeholders, constraints, and intervention points. Submit a 2–3 page individual report.

Teams & Projects.

In Class: Problem discovery and empathize work.

Assignment: Propose a Project — individually, interview real users or study the field, choose an initial project track, and use AI to simulate and critique a user/customer perspective. As a team, share findings, select one project, and submit the project description with rationale.

Module 3 — Judgment Under Evidence — Critical Thinking and the State of the Art

Broader Insight — Mindset & Worldview.

Lecture: The river, the beaver, the software death cross, and the democratization of capability. Critical thinking I: the Paul–Elder framework, hallucination and bias checks, and the principle that plausible is not true. Mindset objective: skill and judgment increasingly matter more than pedigree.

Reading: Ch. 5 “The River of Intelligence and the Beaver's Dam” and Ch. 11 “What the Data Shows.”

Assignment: Fill in the AI Gaps — audit the evidence behind your project concept; identify supported claims, assumptions, and places where AI may have introduced hallucination, bias, or overconfidence. Submit corrected claims and next steps for validation.

Code & Theory — AI Tools & Practice.

Practice: Technology-independent solution research: start from the project requirement, not a preferred product or platform; use AI to support literature review, technical scanning, and rapid comparison.

Assignment: State-of-the-Art Scan — compare at least three approaches to one key project requirement (maturity, evidence, limitations, cost, access, data needs, risk); select and learn one well enough to explain its mechanism or run a basic test. Submit a landscape summary, source list, comparison, and recommendation.

Teams & Projects.

In Class: Use “Start with Story” to draft the team's one-paragraph narrative; use the NABC format (Need, Approach, Benefit, Competition) to structure it.

Assignment: As a team, submit the problem statement and NABC story, noting how the state-of-the-art scan changed or confirmed the approach. Individually, use AI to stress-test the problem statement and record the strongest critique.

Module 4 — The Candle — Authorship, Judgment, and Discomfort

Broader Insight — Mindset & Worldview.

Lecture: The candle and judgment: questions versus answers, what only humans originate, and authorship, accountability, and judgment as scarce assets. Comfort-zone work through a meaningful question normally avoided.

Reading: Ch. 6 “The Candle in the Darkness” and Ch. 7 “Who Is Writing This Book?”

Assignment: The Candle Question — ask a real person a meaningful question you would normally avoid; AI may help prepare possibilities, but you choose and ask the final question yourself. Reflect on the discomfort, the answer, and what it revealed about your judgment.

Code & Theory — AI Tools & Practice.

Practice: “Don't build the shoe; build the factory” — design the project as a repeatable system (AI, automation, conventional software, human process, or a mix chosen to fit the requirement) rather than a one-time deliverable.

Assignment: Factory, Not Shoe — build the simplest possible prototype of the project system, showing input, method/tool, processing step, output, and human checkpoint. Submit the prototype with what worked, what failed, and one risk or limitation.

Teams & Projects.

In Class: Ideate to prototype: develop a presentation covering the problem, approach, user-facing solution, the system design, and comparison with alternatives.

Assignment: Individually, write what changed after feedback and what the next build must prove.

Module 5 — Friction, Flow, and Scope Lock — From Exploration to Execution

Broader Insight — Mindset & Worldview.

Lecture: Friction, flow, and the smooth; grit and deliberate practice; scope lock; trust and psychological safety; the shift from exploration to execution. Mindset objective: let setbacks reveal what the system must handle, then convert ambition into a buildable plan before complexity outruns coordination.

Reading: Ch. 12 “Flow,” Ch. 13 “Friction Has Not Disappeared,” Ch. 15 “The Boulder, the Believer, and the Beaver,” Ch. 19 “The Software Death Cross.”

Assignment: (1) Setbacks and Grit — log what broke, what you tried, what you learned, and what the system must handle next. (2) Mid-course reflection and team charter for the build half.

Code & Theory — AI Tools & Practice.

Practice: Evidence before polish, followed by architecture and scope review: define success, test an early version, and identify the simplest buildable form; review technical, data, reliability, privacy, operational, and human-decision risk.

Assignment: (1) Quality — one success metric, one failure warning sign, and a simple evidence-collection plan. (2) Build Plan — milestones, simplest buildable version, chosen methods/tools with rationale, major risks, human checkpoints, and next iteration.

Teams & Projects.

In Class: Mid-Course Progress and Scope Review — each team demonstrates v0, receives structured peer and mentor feedback, and presents evidence, technology choices, system design, and a proposed scope lock.

Assignment: As a team, turn feedback into a prioritized iteration backlog and update the build plan. Individually, identify the critique you least wanted to hear and name one behavior the team needs from you during the build half.

Module 6 — Stewardship — Risk, Ethics, and Observability

Broader Insight — Mindset & Worldview.

Lecture: The Luddites and the AI pattern: powerful technologies change who has power, who benefits, and who bears the cost. Ask: who might this help, who might it hurt, what could go wrong if it scales? Mindset objective: stewardship — ask “should I?” before “can I?”

Reading: Ch. 8 “The Luddites” and Ch. 17 “The Pattern.”

Assignment: Risks and Ethics — complete a dual-use and stakeholder-risk audit identifying beneficiaries, potential harms, misuse, bias, privacy and security risks, and the human safeguards needed before scaling.

Code & Theory — AI Tools & Practice.

Practice: Workflow and system testing and debugging: run the project on a small test case, make each step observable, and diagnose failure modes (bad input, weak assumptions, missing data, unsuitable technology, poor handoffs, unreliable output, insufficient human review).

Assignment: Observability — create a system log showing input, each human or technology step, output, the weakness observed, and the planned change before the next test.

Teams & Projects.

In Class: Progress Review — present the updated project, evidence, technology choices, system workflow, and remaining backlog to a new peer team.

Assignment: As a team, submit the updated demo plan and remaining backlog. Individually, complete peer review scores and identify one safeguard or design change to add before demonstration.

Module 7 — Attention and Care — Human-Centered Evaluation

Broader Insight — Mindset & Worldview.

Lecture: Attentional ecology and care: attention as a scarce resource, care as the source of quality, and emotional intelligence for users and teammates. Mindset objective: powerful systems still depend on privacy, attention, care, and human relationships to produce humane outcomes.

Reading: Ch. 9 “The Secret Garden” and Ch. 16 “Attentional Ecology.”

Assignment: A short user-care reflection drawn from a real test conversation or teammate interaction.

Code & Theory — AI Tools & Practice.

Practice: Human-centered evaluation: observe how a user provides input, how the project transforms it, and where it helps or fails; identify human-in-the-loop points (who defines the goal, checks the result, decides acceptability, and improves the system).

Assignment: Observation and Human in the Loop — submit an observation memo naming one place the technology helped, one where it failed or overreached, and one where a human must remain responsible, applied to your project.

Teams & Projects.

In Class: Prepare for Demo Day: complete final user testing, incorporate the most important feedback, and confirm alignment among story, evidence, demonstration, and safeguards.

Assignment: None separately assigned; folded into Demo Day preparation.

Module 8 — Signing Your Work — Demo Day and the Foundational Question, Revisited

Broader Insight — Mindset & Worldview.

Lecture: Building, explaining, and signing your work. AI lowers the barrier to creation, but advantage goes to those who create real value and tell a clear story. Revisit the foundational question: what am I for? Mindset objective: move from collecting skills to directing intelligence and other capabilities toward impact.

Reading: Ch. 10 “The Aesthetics of the Smooth,” Ch. 14 “The Democratization of Capability,” Ch. 18 “Leading After the Orange Pill,” Ch. 20 “The Sunrise.”

Assignment: (1) News Story — the user, problem, system, value, evidence, and strongest objection, written as a one-page third-person story. (2) A final answer to the foundational question: what are you for, now that you can direct intelligence?

Code & Theory — AI Tools & Practice.

Practice: Robustness, edge cases, failure modes, reproducibility, and release readiness — the difference between a demonstration and a dependable product or research result.

Assignment: Complete a reliability checklist and release-candidate plan, including methods and tools used, why they were selected, alternatives considered, and an evaluation of strengths and limitations.

Teams & Projects.

In Class: Demo Day and Final Reflection — present the final problem/solution story and a demonstration or video of the system in action.

Assignment: Final presentation and demonstration; updated news story; reflection on learning; list and evaluation of methods and tools used; final peer review and growth report.

ASSESSMENT AND GRADING

The course is graded Pass / No Pass, consistent with its status as research credit. The pass/no-pass determination is based on a curve intended to distinguish students who made genuine accomplishments from those who did not invest effort in the learning experience; students who make a good-faith effort to complete the course will pass.

Grade Allocation

Component	Weight	Points
AI-led, individually customized assignments (homework)	20%	20
Final project + peer review • Creativity and clarity of project intention — 20 pts • Effort and precision of execution — 20 pts	40%	40
Class participation & human-skills exercises	20%	20
Reflection portfolio	20%	20
Total	100%	100

Peer Assessment Process

As part of the final project grade, each student anonymously scores every teammate from 0 to 3 (none / low / medium / high) on three dimensions — Narrative, Design, and Technical Implementation — with a brief written justification for each score. An illustrative example of the resulting summary appears in Appendix A.

COURSE POLICIES

Attendance and Participation

Because this is an applied, team-based project course, attendance and participation are integral to the learning outcomes. Students who are not actively participating in the first weeks of the course may be dropped. Class participation and human-skills exercises constitute approximately 20% of the final grade. Students should not attend class while unwell; sessions are recorded to accommodate excused absences.

Late Work

Late assignments are penalized 50% unless covered by a late day. Each student is allotted seven late days across the term to apply to individual assignments without penalty (any fraction of a day counts as a full day). The final project must be submitted on time, as it is tied to the live Demo Day presentations and the course grading timeline.

Academic Integrity and the Use of Artificial Intelligence

The use of AI in this course is required, not prohibited — but students remain the human in the loop and are expected to sign their own work. Students may use AI for any part of their coursework; they must be able to defend every decision made and disclose how AI contributed to it. Submitting AI output a student cannot personally stand behind, or presenting others' work as one's own, is a violation of academic integrity. The intent of this policy is the opposite of permitting dishonesty: it is to make students better and faster with AI, while preserving the judgment that remains theirs alone.

Inclusion, Accommodations, and Support

The course is committed to a learning environment that welcomes all students and a diversity of perspectives, which the instructors regard as essential to building applications that serve society well. Students affected by

circumstances outside class, or who require accommodations for any reason, are encouraged to speak with the instructors, who are available as a resource; students should also register with their institution's disability services office where applicable [Institution Disability Services Office]. Students should use the name and pronouns they prefer.

STAFF AND OFFICE HOURS

Role	Name & Email	Office Hours
Faculty	Edo Segal — edo@touchcast.com	By appointment and after class
Faculty	Ikhlaq Sidhu — [email]	By appointment and after class
Coach / GSI	[TBD]	[TBD]
AI Twin-Mentors	Edo & Ikhlaq digital twins (companion platform)	Available 24/7 on the platform

APPENDIX A: PEER-REVIEW SUMMARY (ILLUSTRATION)

The table below illustrates the format of a peer-review summary for a hypothetical four-person team, as generated from anonymous teammate scoring (0–3 scale) on Narrative, Design, and Technical Implementation. Actual per-team results are compiled and returned by the instructional team.

Area	Member 2	Member 3	Member 4	Average
Narrative	3.00	2.00	2.00	2.33
Design	0.00	1.00	1.00	0.67
Technical	2.00	2.00	3.00	2.33
Total	5.00	5.00	6.00	5.33