



WEEK 4 CLASS ACTIVITY — “TIER IT” + ASSESSMENT 1 CLINIC

Legal & Regulatory · LO3, LO4 · ~60 minutes (two parts)

Goal. Part 1 — get fluent in the EU AI Act’s four risk tiers by sorting real systems (and your own product) into them. Part 2 — use a peer-review clinic to lift your Assessment 1 before it’s due this week.

Teams of 3–4 (cross-campus via Google Meet). Project deck: [activity_slides.pptx](#).

THE FOUR TIERS (YOUR REFERENCE)

Tier	Rule	Examples
 Unacceptable	banned (Article 5)	govt social scoring; manipulative AI; real-time face recognition for policing
 High risk	heavily regulated (Annex III)	credit scoring, hiring, education marking, biometrics, essential services
 Limited risk	just disclose it’s AI (Article 50)	chatbots; AI-generated content; deepfakes
 Minimal risk	almost no rules	spam filters, game AI, recommendations, forecasting

PART 1 — “TIER IT” (≈35 MIN)

Step 1 — Sort these eight (≈18 min)

As a team, put each system in a tier. Then **argue the close calls** — which ones jump from “limited” to “high”, and why?



	System	Tier	Why
A	A government system that ranks citizens by their behaviour		
B	A bank's credit-scoring model that approves or rejects loans		
C	A tool that ranks job applicants' CVs for hiring		
D	Real-time facial recognition for policing in public spaces		
E	A customer-service chatbot on a retail website		
F	An AI that writes marketing captions and ad images		
G	An email spam filter		
H	A video game's enemy-character AI		

Step 2 — Tier your own product (≈17 min)

Place your **Assessment 1 product** in a tier, and write the **obligations** that follow:

- **Minimal** → no specific AI Act duties (just existing privacy/consumer law).
- **Limited** → disclose it's AI / that content is AI-generated.
- **High** → heavy duties → consider **redesigning to step down a tier** (make the AI *assist*, not *decide*).

And whatever your tier: **privacy law (PDPA) still applies**. Note what personal data you hold, whose, and where it goes.

PART 2 — ASSESSMENT 1 CLINIC (≈25 MIN)

Swap your near-complete draft with another student. Review theirs against the rubric — be the tough, kind reviewer you'd want. Tick or flag each:

- ☐ **Problem** — specific & real (not “improve X”)?
- ☐ **AI landscape** — ≥3 options across stack layers, with real examples?
- ☐ **Justified choice** — uses the Week 1 decision aid (not just the fanciest option)?
- ☐ **Risks & mitigations** — real risks (Week 3 scorecard) + concrete fixes (the three moves)?
- ☐ **BMC** — all nine blocks, and do they *agree* with each other?
- ☐ **AI-use** — AI passages annotated; specific, non-generic?

Give your partner **one strength** and the **one change** that would most lift their grade. Then fold the feedback in and **submit on time**.



FACILITATOR NOTES (DO NOT DISTRIBUTE)

- **“Tier It” answer sketch:** A = ☐ prohibited (social scoring); B = ☒ high (credit scoring, essential services); C = ☒ high (employment); D = ☐ prohibited (real-time biometric policing); E = ☒ limited (disclose the bot); F = ☒ limited (disclose AI content); G = ☐ minimal; H = ☐ minimal. The teaching points: (1) the *same* technology can be different tiers by *use* — face recognition for unlocking your phone (minimal) vs real-time policing (prohibited); (2) credit/hiring/education are the high-risk traps a student fintech/edtech idea can stumble into; (3) the fix is usually to **design down** — assist, don’t decide.
- **Part 1 facilitation:** push the close calls — “what would move E from limited to high?” (if the chatbot started making binding decisions about essential services). Keep it to ~18 min so the clinic gets its time.
- **Part 2 is the priority** — Asst 1 is due this week. Protect the 25 minutes. Teaching team circulates as a second reviewer for any student without a partner.
- **Cross-campus:** mixed teams sort one shared tier-board (Jamboard); for the clinic, pair within or across campuses via Meet.
- Close by reminding everyone of the submission deadline and the Moodle link.